Published in final edited form as:

J Am Stat Assoc. 2017 ; 112(517): 169–187. doi:10.1080/01621459.2015.1093947.

Residual Weighted Learning for Estimating Individualized Treatment Rules

Xin Zhou [PhD student],

Department of Biostatistics, University of North Carolina at Chapel Hill, NC 27599

(xinzhou@live.unc.edu)

Nicole Mayer-Hamblett [Associate Professor] [Adjunct Associate Professor],

Department of Pediatrics and Department of Biostatistics, University of Washington

(nicole.hamblett@seattlechildrens.org)

Umer Khan [statistician], and

Cystic Fibrosis Foundation Therapeutics Development Network, Seattle Childrens Hospital

(umer.khan@seattlechildrens.org)

Michael R. Kosorok [W. R. Kenan, Jr. Distinguished Professor and Chair]

Department of Biostatistics, and Professor, Department of Statistics and Operations Research,
University of North Carolina at Chapel Hill, Chapel Hill, NC 27599 (kosorok@unc.edu)

Abstract

Personalized medicine has received increasing attention among statisticians, computer scientists, and clinical practitioners. A major component of personalized medicine is the estimation of individualized treatment rules (ITRs). Recently, Zhao et al. (2012) proposed outcome weighted learning (OWL) to construct ITRs that directly optimize the clinical outcome. Although OWL opens the door to introducing machine learning techniques to optimal treatment regimes, it still has some problems in performance. (1) The estimated ITR of OWL is affected by a simple shift of the outcome. (2) The rule from OWL tries to keep treatment assignments that subjects actually received. (3) There is no variable selection mechanism with OWL. All of them weaken the finite sample performance of OWL. In this article, we propose a general framework, called Residual Weighted Learning (RWL), to alleviate these problems, and hence to improve finite sample performance. Unlike OWL which weights misclassification errors by clinical outcomes, RWL weights these errors by residuals of the outcome from a regression fit on clinical covariates excluding treatment assignment. We utilize the smoothed ramp loss function in RWL, and provide a difference of convex (d.c.) algorithm to solve the corresponding non-convex optimization problem. By estimating residuals with linear models or generalized linear models, RWL can effectively deal with different types of outcomes, such as continuous, binary and count outcomes. We also propose variable selection methods for linear and nonlinear rules, respectively, to further improve the performance. We show that the resulting estimator of the treatment rule is consistent. We further obtain a rate of convergence for the difference between the expected outcome using the estimated ITR and that of the optimal treatment rule. The performance of the proposed RWL methods is illustrated in simulation studies and in an analysis of cystic fibrosis clinical trial data.

Keywords

Optimal Treatment Regime; RKHS; Universal consistency; Convergence Rate; Residuals; Martingale residuals

1 Introduction

Personalized medicine is a medical paradigm that utilizes individual patient information to optimize patients health care. Recently, personalized medicine has received much attention among statisticians, computer scientists, and clinical practitioners. The primary motivation is the well-established fact that patients often show significant heterogeneity in response to treatments. For instance, in a recent study, it was demonstrated that the optimal timing for the initiation of antiretroviral therapy (ART) varies in patients co-infected with human immunodeficiency virus and tuberculosis. Patients with CD4+ T-cell counts of less than 50 per cubic millimeter benefited substantially from earlier ART with a lower rate of new AIDS-defining illnesses and mortality as compared with later ART, while those with larger CD4+ T-cell counts did not have such a benefit (Havlir et al. 2011). The inherent heterogeneity across patients suggests a transition from the traditional “one size fits all” approach to modern personalized medicine.

A major component of personalized medicine is the estimation of individualized treatment rules (ITRs). Formally, the goal is to seek a rule that assigns a treatment, from among a set of possible treatments, to a patient based on his or her clinical, prognostic or genomic characteristics. The individualized treatment rules are also called optimal treatment regimes. There is a significant literature on individualized treatment strategies based on data from clinical trials or observational studies (Murphy 2003; Robins 2004; Zhang et al. 2012; Zhao et al. 2009). Much of the work relies on modeling either the conditional mean outcomes or contrasts between mean outcomes. These methods obtain ITRs indirectly by inverting the regression estimates instead of directly optimizing the decision rule. For instance, Qian and Murphy (2011) applied a two-step procedure that first estimates a conditional mean for the outcome and then determines the treatment rule by comparing the conditional means across various treatments. The success of these indirect approaches highly depends on correct specification of posited models and on the precision of model estimates.

In contrast, Zhao et al. (2012) proposed outcome weighted learning (OWL), using data from a randomized clinical trial, to construct an ITR that directly optimizes the clinical outcome. They cast the treatment selection problem as a weighted classification problem, and apply state-of-the-art support vector machines for implementation. This approach opens the door to introducing machine learning techniques into this area. However, there is still significant room for improved performance. First, the estimated ITR of OWL is affected by a simple shift of the outcome. This behavior makes the estimate of OWL unstable. Second, since OWL requires the outcome to be nonnegative, OWL works similarly to weighted classification, in which misclassification errors, the differences between the estimated and true treatment assignments, are targeted to be reduced. Hence the ITR estimated by OWL tries to keep treatment assignments that subjects actually received. This is not always ideal

since treatments are randomly assigned in the trial, and the probability is slim that the majority of subjects are assigned optimal treatments. Third, OWL does not have variable selection features. When there are many clinical covariates which are not related to the heterogeneous treatment effects, variable selection is critical to the performance.

To alleviate these problems, we propose a new method, called Residual Weighted Learning (RWL). Unlike OWL which weights misclassification errors by clinical outcomes, RWL weights these errors by residuals from a regression fit of the outcome. The predictors of the regression model include clinical covariates, but exclude treatment assignment. Thus the residuals better reflect the heterogeneity of treatment effects. Since some residuals are negative, the hinge loss function used in OWL is not appropriate. We instead utilize the smoothed ramp loss function, and provide a difference of convex (d.c.) algorithm to solve the corresponding non-convex optimization problem. The smoothed ramp loss resembles the ramp loss (Collobert et al. 2006), but it is smooth everywhere. It is well known that ramp loss related methods are shown to be robust to outliers (Wu and Liu 2007). The robustness to outliers for the smoothed ramp loss is helpful for RWL, especially when residuals are poorly estimated. Moreover, through using residuals, RWL is able to deal relatively easily with almost all types of outcomes. For example, RWL can work with generalized linear models to construct ITRs for count/rate outcomes. We also propose variable selection approaches in RWL for linear and nonlinear rules, respectively, to further improve finite sample performance.

The theoretical analysis of RWL focuses on two aspects, universal consistency and convergence rate. The notion of universal consistency is borrowed from machine learning. It requires for a learning method that when the sample size approaches infinity the method eventually learns the Bayes rule without knowing any specifics of the distribution of the data. We show that RWL with a universal kernel (e.g. Gaussian RBF kernel) is universally consistent. In machine learning, there is a famous “no-free-lunch theorem”, which states that the convergence rate of any particular learning rule may be arbitrarily slow (Devroye et al. 1996). In this article, we prove the “no-free-lunch theorem” for finding ITRs. Thus the rate of convergence studies for a particular rule must necessarily be accompanied by conditions on the distribution of the data. For RWL with Gaussian RBF kernel, we show that under the geometric noise condition (Steinwart and Scovel 2007) the convergence rate is as high as $n^{-1/3}$.

At first glance, one may think that there is not a large difference between OWL and RWL except that RWL uses residuals as alternative outcomes. Actually, RWL enjoys many benefits from this simple modification. First, by using residuals of outcomes, RWL is able to reduce the variability introduced by the original outcomes. Second, since the numbers of subjects with positive and negative residuals are generally balanced, the ITR determined by RWL favors neither the treatment assignments that subjects actually received nor their opposites. Because of the above reasons, RWL improves finite sample performance. Third, RWL possesses location-scale invariance with respect to the original outcomes. Specifically, the estimated rule of RWL is invariant to a shift of the outcome; it is invariant to a scaling of the outcome with a positive number; the rule from RWL that maximizes the outcome is opposite to the rule that minimizes the outcome. These are intuitively sensible.

The contributions of this article are summarized as follows. (1) We propose the general framework of Residual Weighted Learning to estimate individualized treatment rules. By estimating residuals with linear or generalized linear models, RWL can effectively deal with different types of outcomes, such as continuous, binary and count outcomes. For censored survival outcomes, RWL could potentially utilize martingale residuals, although theoretical justification is still under development. (2) We develop variable selection techniques in RWL to further improve performance. (3) We present a comprehensive theoretical analysis of RWL on universal consistency and convergence rate. (4) As a by-product, we show the “no-free-lunch theorem” for ITRs that the convergence rate of any particular rule may be arbitrarily slow. This is a generic result, and it applies to any algorithm for optimal treatment regimes.

The remainder of the article is organized as follows. In Section 2, we discuss Outcome Weighted Learning (OWL) and propose Residual Weighted Learning (RWL) to improve finite sample performance for continuous outcomes. Then we develop a general framework for RWL to handle other types of outcomes, and use binary and count/rate outcomes as examples. In Section 3, we establish consistency and convergence rate results for the estimated rules. The variable selection techniques for RWL are discussed in Section 4. We present simulation studies to evaluate performance of the proposed methods in Section 5. The method is then illustrated on the EPIC cystic fibrosis randomized clinical trial (Treggiari et al. 2009, 2011) in Section 6. We conclude the article with a discussion in Section 7. All proofs are given in the supplementary material.

2 Methodology

2.1 Outcome Weighted Learning

Consider a two-arm randomized trial. We observe a triplet $(\mathbf{X}, A, R)$ from each patient, where $\mathbf{X} = (X_1, \dots, X_p)^T \in \mathcal{X}$ denotes the patient's clinical covariates, $A \in \mathcal{A} = \{1, -1\}$ denotes the treatment assignment, and R is the observed clinical outcome, also called the “reward” in the literature on reinforcement learning. We assume that R is bounded, and larger values of R are more desirable. An individualized treatment rule (ITR) is a function from $\mathcal{X}$ to $\mathcal{A}$. Let $\pi(a, \mathbf{x}) := P(A = a | \mathbf{X} = \mathbf{x})$ be the probability of being assigned treatment a for patients with clinical covariates $\mathbf{x}$. It is predefined in the trial design. Here we consider a general situation. In most clinical trials, the treatment assignment is independent of $\mathbf{X}$, but in some designs, such as stratified designs, A may depend on $\mathbf{X}$. We assume $\pi(a, \mathbf{x}) > 0$ for all $a \in \mathcal{A}$ and $\mathbf{x} \in \mathcal{X}$.

An optimal ITR is a rule that maximizes the expected outcome under this rule. Mathematically, the expected outcome under any ITR d is given as

$$\mathbb{E}(R | A = d(\mathbf{X})) = \mathbb{E} \left(\frac{R}{\pi(A, \mathbf{X})} \mathbb{I}(A = d(\mathbf{X})) \right), \quad (1)$$

where $\mathbb{I}(\cdot)$ is the indicator function. Interested readers may refer to Qian and Murphy (2011) and Zhao et al. (2012) regarding the derivation of (1). This expectation is called the value

function associated with the rule d , and is denoted $V(d)$. In other words, the value function $V(d)$ is the expected value of R given that the rule $d(X)$ is applied to the given population of patients. An optimal ITR d^* is a rule that maximizes $V(d)$. Finding d^* is equivalent to the following minimization problem:

$$d^* \in \arg \min_d \mathbb{E} \left(\frac{R}{\pi(A, \mathbf{X})} \mathbb{I}(A \neq d(\mathbf{X})) \right). \quad (2)$$

Zhao et al. (2012) viewed this as a weighted classification problem, in which one wants to classify A using X but also weights each misclassification error by $R/\pi(A, X)$.

Assume that the observed data $\{(\mathbf{x}_i, a_i, r_i) : i = 1, \dots, n\}$ are collected independently. For any decision function $f(\mathbf{x})$, let $d_f(\mathbf{x}) = \text{sign}(f(\mathbf{x}))$ be the associated rule, where $\text{sign}(u) = 1$ for $u > 0$ and -1 otherwise. The particular choice of the value of $\text{sign}(0)$ is not important. In this article, we fix $\text{sign}(0) = -1$. Using the observed data, the weighted classification error in (2) can be approximated by the empirical risk

$$\frac{1}{n} \sum_{i=1}^n \frac{r_i}{\pi(a_i, \mathbf{x}_i)} \mathbb{I}(d_f(\mathbf{x}_i) \neq a_i). \quad (3)$$

The optimal decision function $\hat{f}(\mathbf{x})$ is the one that minimizes the outcome weighted error (3).

It is well known that empirical risk minimization for a classification problem with the 0-1 loss function is an NP-hard problem. To alleviate this difficulty, one often finds a surrogate loss to replace the 0-1 loss. Outcome weighted learning (OWL) proposed by Zhao et al. (2012) uses the hinge loss function, and also applies the regularization technique used in the support vector machine (SVM) (Vapnik 1998). In other words, instead of minimizing (3), OWL aims to minimize

$$\frac{1}{n} \sum_{i=1}^n \frac{r_i}{\pi(a_i, \mathbf{x}_i)} (1 - a_i f(\mathbf{x}_i))_+ + \lambda \|f\|^2, \quad (4)$$

where $(u)_+ = \max(u, 0)$ is the positive part of u , $\|f\|$ is some norm for f , and λ is a tuning parameter controlling the trade-off between empirical risk and complexity of the decision function f .

OWL opens the door to the application of statistical learning techniques to personalized medicine. However, this approach is not perfect. First, a simple shift on the outcome R should not affect the optimal decision rule, as seen from (1) and (2). That is, d^* in (2) does not change if R is replaced by $R + c$, for any constant c . Unfortunately this nice invariance property does not hold for the decision function $\hat{f}(\mathbf{x})$ of OWL in (4) when r_i is replaced by $r_i + c$. Consider an extreme case where c is very large and $\pi(1, \mathbf{x}) = \pi(-1, \mathbf{x}) = 0.5$ for all

$x \in \mathcal{X}$, the weights $(r_i + c)/\pi(a_i, x_i)$ are almost identical, and the weighted problem is approximately transformed to an unweighted one. Hence the performance of OWL can be affected by a simple shift of R . Second, OWL further assumes that R is nonnegative to gain computational efficiency from convex programming. A direct consequence of this assumption, as seen from (3), is that the treatment regime $d_f(x_i)$ tends to match the treatment a_i that was actually assigned to the patient, especially when the decision function f is chosen from a rich class of functions. This property is not ideal for data from a randomized clinical trial, since treatments are actually randomly assigned to patients.

2.2 Residual Weighted Learning

In this section, we only consider the case that the outcome R is continuous, and extend the framework to other types of outcomes in Section 2.4.

As demonstrated previously, the decision rule in (2) is invariant to a shift of outcome R by any constant. Moreover, Lemma 2.1 shows that d^* in (2) is invariant under a shift of R by a function of $\mathbf{X}$. That is, d^* remains unchanged if R is replaced by $R - g(\mathbf{X})$ for any function g , as long as g is not related to d . So there is an optimal solution $\hat{f}$ in (4) corresponding to a particular function g , when r_i is replaced by $r_i - g(x_i)$, and the associated rule $\hat{d}_f = \text{sign}(\hat{f})$ can be seen as an estimated optimal rule in (2). As g varies, we obtain a collection of $\hat{d}_f$'s.

When the sample size is very large, these $\hat{d}_f$'s perform similarly by the infinite sample property of OWL shown in Zhao et al. (2012). However, when the sample size is limited, as in clinical settings, the choice of g is critical to the performance of OWL. Zhao et al. (2012) did not delve deeply into this problem. In this article, we will provide a solution to this, and improve finite sample performance.

Lemma 2.1. *For any measurable $g : \mathcal{X} \rightarrow \mathbb{R}$ and any probability distribution for $(\mathbf{X}, A, R)$,*

$$\mathbb{E} \left(\frac{R - g(\mathbf{X})}{\pi(A, \mathbf{X})} \mathbb{I}(A \neq d(\mathbf{X})) \right) = \mathbb{E} \left(\frac{R}{\pi(A, \mathbf{X})} \mathbb{I}(A \neq d(\mathbf{X})) \right) - \mathbb{E}(g(\mathbf{X})).$$

Intuitively, the function g with the smallest variance of $\frac{R - g(\mathbf{X})}{\pi(A, \mathbf{X})} \mathbb{I}(A \neq d(\mathbf{X}))$ is a good choice. However, such a function depends on the decision rule d , as shown in the following theorem. The proof of the theorem is provided in the supplementary material.

Theorem 2.2. *Among all measurable*

$g : \mathcal{X} \rightarrow \mathbb{R}$, $\tilde{g}(\mathbf{X}) = \mathbb{E} \left(\frac{R}{\pi(A, \mathbf{X})} \mathbb{I}(A \neq d(\mathbf{X})) \mid \mathbf{X} \right)$ *is the function g that minimizes the variance of*

$$\frac{R - g(\mathbf{X})}{\pi(A, \mathbf{X})} \mathbb{I}(A \neq d(\mathbf{X})).$$

Our purpose is to find a function g , which is not related to d , to reduce the variance of

$\frac{R - g(\mathbf{X})}{\pi(A, \mathbf{X})} \mathbb{I}(A \neq d(\mathbf{X}))$ as much as possible. As shown in the proof, the minimizer $\tilde{g}$ can be written as,

$$\tilde{g}(\mathbf{X}) = \mathbb{E}(R|\mathbf{X}, A=1) \mathbb{I}(d(\mathbf{X}) \neq 1) + \mathbb{E}(R|\mathbf{X}, A=-1) \mathbb{I}(d(\mathbf{X}) \neq -1).$$

That is, $\tilde{g}(\mathbf{X})$ jumps between $\mathbb{E}(R|\mathbf{X}, A=1)$ and $\mathbb{E}(R|\mathbf{X}, A=-1)$ as $d(\mathbf{X})$ varies. Hence when d is unknown, a reasonable choice of g is

$$g^*(\mathbf{X}) = \frac{\mathbb{E}(R|\mathbf{X}, A=1) + \mathbb{E}(R|\mathbf{X}, A=-1)}{2} = \mathbb{E}\left(\frac{R}{2\pi(A, \mathbf{X})} | \mathbf{X}\right). \quad (5)$$

We propose to minimize the following empirical risk, rather than the original one in (3):

$$\frac{1}{n} \sum_{i=1}^n \frac{r_i - \hat{g}^*(\mathbf{x}_i)}{\pi(a_i, \mathbf{x}_i)} \mathbb{I}(d_f(\mathbf{x}_i) \neq a_i), \quad (6)$$

where $\hat{g}^*$ is an estimate of g^* . For simplicity, let $\hat{r}_i = r_i - \hat{g}^*(\mathbf{x}_i)$ be the estimated residual. Here, we do not weight misclassification errors by clinical outcomes as OWL does, and instead we weight them by residuals from a regression fit of outcomes. Thus the optimal decision function is invariant under any translation of clinical outcomes. We call the method Residual Weighted Learning (RWL).

RWL also has the following interpretation. The outcome R can be characterized as

$$R = \mu(\mathbf{X}) + \delta(\mathbf{X}) \cdot A + \epsilon,$$

where ϵ is the mean zero random error. $\mu(\mathbf{X})$ reflects common effects of clinical covariates $\mathbf{X}$ for both treatment arms. It is easy to see that

$\mu(\mathbf{X}) = (\mathbb{E}(R|\mathbf{X}, A=1) + \mathbb{E}(R|\mathbf{X}, A=-1)) / 2$, so the residual, $R - g^*(\mathbf{X})$, in RWL is just $\delta(\mathbf{X}) \cdot A + \epsilon$, which captures all sources of heterogeneous treatment effects. The rationale of optimal treatment regimes is to keep treatment assignments that subjects have actually received if those subjects are observed to have large outcomes, and to switch assignments if outcomes are small. However, the largeness and smallness for outcomes are relative. A rather large outcome may still be considered as small when compared with subjects having similar clinical covariates, as shown in Figure 1. Outcomes are not comparable among subjects with different clinical covariates, while the residual, by removing common covariates effects, is a better measurement. Figure 1 illustrates how residuals work using an example with a single covariate X . The raw data are shown on the left, and residuals are shown on the right. Residuals are comparable among subjects, and larger residuals represent better outcomes.

Another benefit from using residuals is a clear cut-off, *i.e.* 0, to distinguish between subjects with good and poor clinical outcomes. To minimize the empirical risk in (6), for subjects with positive residuals, RWL is apt to recommend the same treatment assignments that subjects have actually received; for subjects with negative residuals, RWL is more likely to give the opposite treatment assignments to what they have received. The optimal ITR is the one maximizing the conditional expected outcome given in (1). An empirical estimator of (1) is,

$$\frac{1}{n} \sum_{i=1}^n \frac{r_i}{\pi(a_i, \mathbf{x}_i)} \mathbb{I}(d(\mathbf{x}_i) = a_i). \quad (7)$$

Though it is unbiased, it may give an estimate outside the range of R . While for finite samples a better estimator of (1) as shown in Murphy (2005) is,

$$\frac{\frac{1}{n} \sum_{i=1}^n \frac{r_i \mathbb{I}(d(\mathbf{x}_i) = a_i)}{\pi(a_i, \mathbf{x}_i)}}{\frac{1}{n} \sum_{i=1}^n \frac{\mathbb{I}(d(\mathbf{x}_i) = a_i)}{\pi(a_i, \mathbf{x}_i)}}. \quad (8)$$

The denominator is an estimator of $\mathbb{E} \left(\frac{\mathbb{I}(A=d(\mathbf{X}))}{\pi(A, \mathbf{X})} \right)$. Similar reasoning as that used in the proof of Lemma 2.1 yields $\mathbb{E} \left(\frac{\mathbb{I}(A=d(\mathbf{X}))}{\pi(A, \mathbf{X})} \right) = 1$. The estimator $\frac{1}{n} \sum_{i=1}^n \frac{\mathbb{I}(d(\mathbf{x}_i) = a_i)}{\pi(a_i, \mathbf{x}_i)}$ is called the treatment matching factor in this article. The factor varies between 0 and 2. If a rule favors treatment assignments that subjects have actually received, its treatment matching factor is greater than 1, while in contrast if a rule prefers the opposite treatment assignments, its treatment matching factor is less than 1. For a randomized clinical trial, we expect that the estimated rule is associated with a treatment matching factor close to 1. For OWL, since all the weights are nonnegative, the estimated rule of OWL tends to keep, if possible, the treatments that subjects actually received. Thus the associated treatment matching factor would be greater than 1, especially when the sample size is small, or when a complicated rule is applied. Hence the estimator in (8) might not be large even though (7) is maximized. RWL alleviates this problem by using residuals to make an initial guess on the optimal rule. Owing to the balance between subjects with positive and negative residuals, RWL implicitly finds a rule with its treatment matching factor close to 1.

There are many ways to estimate g^* . In this article, we consider two models. The first one is the main effects model. Assume that $\mathbb{E} \left(\frac{R}{2\pi(A, \mathbf{X})} | \mathbf{X} \right) = \beta_0 + \mathbf{X}^T \boldsymbol{\beta}$, where $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^T$. The estimates $\hat{\beta}$ and $\hat{\beta}_0$ can be obtained by minimizing the sum of weighted squares,

$$\sum_{i=1}^n \frac{1}{2\pi(a_i, \mathbf{x}_i)} (r_i - \beta_0 - \mathbf{x}_i^T \boldsymbol{\beta})^2.$$

It can be solved easily by almost any statistical software. Then the estimate is

$\hat{g}^*(x) = \hat{\beta}_0 + x^T \hat{\beta}$. The second is the null model. Assume that $\mathbb{E}\left(\frac{R}{2\pi(A, \mathbf{X})} | \mathbf{X}\right) = \beta_0$. It is easy to obtain $\hat{g}^*(x) = \hat{\beta}_0 = \frac{\sum_{i=1}^n r_i}{\sum_{i=1}^n \pi(a_i, x_i)} / \frac{1}{\sum_{i=1}^n \pi(a_i, x_i)}$.

In summary, we propose a method called Residual Weighted Learning to identify the optimal ITR by minimizing the residual weighted classification error (6). The impact of using residuals is two-fold: it stabilizes the variance of the value function and controls the treatment matching factor. Both improve finite sample performance.

2.3 Implementation of RWL

As in OWL, we intend to use a surrogate loss function to replace the 0-1 loss in (6). Since some residuals are negative, convex surrogate functions do not work here. We consider a non-convex loss

$$T(u) = \begin{cases} 0 & \text{if } u \geq 1, \\ (1-u)^2 & \text{if } 0 \leq u < 1, \\ 2 - (1+u)^2 & \text{if } -1 \leq u < 0, \\ 2 & \text{if } u < -1. \end{cases}$$

It is called the smoothed ramp loss in this article. Figure 2 shows the hinge loss, ramp loss and smoothed ramp loss functions. The hinge loss is the loss function used in support vector machines (Vapnik 1998). The ramp loss (Collobert et al. 2006) is also called the truncated hinge loss function (Wu and Liu 2007). It is well known that by truncating the unbounded hinge loss, ramp loss related methods are shown to be robust to outliers in the training data for the classification problem (Wu and Liu 2007). Compared with the ramp loss, the smoothed ramp loss is smooth everywhere. Hence it has computational advantages in optimization. Moreover, the smoothed ramp loss, which resembles the ramp loss, is robust to outliers too. In the framework of RWL, subjects who did not receive optimal treatment assignments but had large positive residuals, or those who did receive optimal assignments but had large negative residuals may be considered as outliers. The robustness to outliers for the smoothed ramp loss is helpful in dealing with outliers in the setting of optimal treatment regimes, especially when residuals are poorly estimated, or when the outcome R has a large variance.

We incorporate RWL into the regularization framework, and aim to minimize

$$\frac{1}{n} \sum_{i=1}^n \frac{\hat{r}_i}{\pi(a_i, x_i)} T(a_i f(x_i)) + \frac{\lambda}{2} \|f\|^2, \quad (9)$$

where $\|f\|$ is some norm for f , and λ is a tuning parameter. Recall that $\hat{r}_i$ is the estimated residual. The smoothed ramp loss function $T(u)$ is symmetric about the point (0; 1) as shown in Figure 2(c). A nice property that comes from the symmetry is that the rule that minimizes

the outcome R (i.e. maximizes $-R$) is just opposite to the rule that maximizes R . This is intuitively sensible. However, OWL does not possess this property.

We derive an algorithm for linear RWL in Section 2.3.1, and then generalize it to the case of nonlinear learning through kernel mapping in Section 2.3.2.

2.3.1 Linear Decision Rule for Optimal ITR—Suppose that the decision function $f(\mathbf{x})$ that minimizes (9) is a linear function of $\mathbf{x}$, i.e. $f(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + b$. Then the associated ITR will assign a subject with clinical covariates $\mathbf{x}$ into treatment 1 if $\mathbf{w}^T \mathbf{x} + b > 0$ and -1 otherwise. In (9), we define $\|f\|$ as the Euclidean norm of $\mathbf{w}$. Then minimizing (9) can be rewritten as

$$\min_{\mathbf{w}, b} \frac{\lambda}{2} \mathbf{w}^T \mathbf{w} + \frac{1}{n} \sum_{i=1}^n \frac{\hat{r}_i}{\pi(a_i, \mathbf{x}_i)} T(a_i (\mathbf{w}^T \mathbf{x}_i + b)). \quad (10)$$

The smoothed ramp loss is a non-convex loss, and as a result, the optimization problem in (10) involves non-convex minimization. This optimization problem is difficult since there are many local minima or stationary points. For instance, any $(\mathbf{w}, b)$ with $\mathbf{w} = \mathbf{0}$ and $|b| = 1$ is a stationary point. Similar with the robust truncated hinge loss support vector machine (Wu and Liu 2007), we apply the d.c. (Difference of Convex) algorithm (An and Tao 1997) to solve this non-convex minimization problem. The d.c. algorithm is also known as the Concave-Convex Procedure (CCCP) in the machine learning community (Yuille and Rangarajan 2003). Assume that an objective function can be rewritten as the sum of a convex part $Q_{\text{vex}}(\Theta)$ and a concave part $Q_{\text{cav}}(\Theta)$. The d.c. algorithm as shown in Algorithm 1 solves the non-convex optimization problem by minimizing a sequence of convex subproblems. One can easily see that the d.c. algorithm is a special case of the Majorize-Minimization (MM) algorithm.

Algorithm 1: The d.c. algorithm for minimizing $Q(\Theta) = Q_{\text{vex}}(\Theta) + Q_{\text{cav}}(\Theta)$.

Initialize $\Theta^{(0)}$

repeat

$$\Theta^{(t+1)} = \operatorname{argmin}_{\Theta} Q_{\text{vex}}(\Theta) + \langle Q'_{\text{cav}}(\Theta^{(t)}), \Theta \rangle$$

until convergence of $\Theta^{(t)}$

Let

$$\phi_s(u) = \begin{cases} 0 & \text{if } u \geq s \\ (s-u)^2 & \text{if } s-1 \leq u < s \\ 2s-2u-1 & \text{if } u < s-1 \end{cases}.$$

Note that ϕ_s is smooth. We have a difference-of-convex decomposition of the smoothed ramp loss,

$$T(u) = \phi_1(u) - \phi_0(u), \quad (11)$$

as shown in Figure 2(d). Denote Θ as $(\mathbf{w}, \mathbf{b})$. Applying (11), the objective function in (10) can be decomposed as

$$Q^s(\Theta) = \underbrace{\frac{\lambda}{2} \mathbf{w}^T \mathbf{w} + \frac{1}{n} \sum_{i=1}^n [\phi_1(u_i) \mathbb{I}(\hat{r}_i > 0) + \phi_0(u_i) \mathbb{I}(\hat{r}_i < 0)]}_{Q_{\text{vec}}(\Theta)} \frac{|\hat{r}_i|}{\pi(a_i, \mathbf{x}_i)} - \underbrace{\frac{1}{n} \sum_{i=1}^n [\phi_1(u_i) \mathbb{I}(\hat{r}_i < 0) + \phi_0(u_i) \mathbb{I}(\hat{r}_i > 0)]}_{Q_{\text{cav}}(\Theta)} \frac{|\hat{r}_i|}{\pi(a_i, \mathbf{x}_i)},$$

where $u_i = a_i(\mathbf{w}^T \mathbf{x}_i + \mathbf{b})$. For simplicity, we introduce the notation,

$$\beta_i = \frac{\partial Q_{\text{cav}}}{\partial u_i} = -\frac{1}{n} \left[\frac{d\phi_1(u_i)}{du_i} \mathbb{I}(\hat{r}_i < 0) + \frac{d\phi_0(u_i)}{du_i} \mathbb{I}(\hat{r}_i > 0) \right] \frac{|\hat{r}_i|}{\pi(a_i, \mathbf{x}_i)}, \quad (12)$$

for $i = 1, \dots, n$. Thus the convex subproblem at the $(t+1)$ 'th iteration of the d.c. algorithm is

$$\min_{\mathbf{w}, \mathbf{b}} \frac{\lambda}{2} \mathbf{w}^T \mathbf{w} + \frac{1}{n} \sum_{i=1}^n [\phi_1(u_i) \mathbb{I}(\hat{r}_i > 0) + \phi_0(u_i) \mathbb{I}(\hat{r}_i < 0)] \frac{|\hat{r}_i|}{\pi(a_i, \mathbf{x}_i)} + \sum_{i=1}^n \beta_i^{(t)} u_i. \quad (13)$$

There are many efficient methods available for solving smooth unconstrained optimization problems. In this article, we use the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) algorithm (Nocedal 1980). L-BFGS is a quasi-Newton method that approximates the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm using a limited amount of computer memory. For more information on the quasi-Newton method and L-BFGS, see Nocedal and Wright (2006).

The procedure is summarized in the following algorithm:

Algorithm 2: The d.c. algorithm for linear RWL with the smoothed ramp loss

Set ϵ to a small quantity, say, 10^{-8}

Initialize $\beta_i^{(0)} = \frac{2|\hat{r}_i|}{n\pi(a_i, \mathbf{x}_i)} \mathbb{I}(\hat{r}_i < 0)$

repeat

- Compute $(\hat{\mathbf{w}}, \hat{b})$ by solving (13),
- Update $u_i = a_i(\hat{\mathbf{w}}^T \mathbf{x}_i + \hat{b})$
- Update $\beta_i^{(t+1)}$ by (12)

until $\|\beta^{(t+1)} - \beta^{(t)}\|_\infty < \epsilon$

2.3.2 Nonlinear Decision rule for Optimal ITR—For a nonlinear decision rule, the decision function $f(\mathbf{x})$ is represented by $h(\mathbf{x}) + b$ with $h(\mathbf{x}) \in \mathcal{H}_K$ and $b \in \mathbb{R}$, where $\mathcal{H}_K$ is a reproducing kernel Hilbert space (RKHS) associated with a Mercer kernel function K . The kernel function $K(\cdot, \cdot)$ is a positive definite function mapping from $\mathcal{X} \times \mathcal{X}$ to $\mathbb{R}$. The norm in $\mathcal{H}_K$, denoted by $\|\cdot\|_K$, is induced by the following inner product:

$$\langle f, g \rangle_K = \sum_{i=1}^n \sum_{j=1}^m \alpha_i \beta_j K(\mathbf{x}_i, \mathbf{x}_j),$$

for $f(\cdot) = \sum_{i=1}^n \alpha_i K(\cdot, \mathbf{x}_i)$ and $g(\cdot) = \sum_{j=1}^m \beta_j K(\cdot, \mathbf{x}_j)$. Then minimizing (9) can be rewritten as

$$\min_{h, b} \frac{\lambda}{2} \|h\|_K^2 + \frac{1}{n} \sum_{i=1}^n \frac{\hat{r}_i}{\pi(a_i, \mathbf{x}_i)} T(a_i(h(\mathbf{x}_i) + b)). \quad (14)$$

Due to the representer theorem (Kimeldorf and Wahba 1971), the nonlinear problem can be reduced to finding finite-dimensional coefficients v_j and $h(\mathbf{x})$ can be represented as

$\sum_{j=1}^n v_j K(\mathbf{x}, \mathbf{x}_j)$. Thus the problem (14) is transformed to

$$\min_{v, b} \frac{\lambda}{2} \sum_{i, j=1}^n v_i v_j K(\mathbf{x}_i, \mathbf{x}_j) + \frac{1}{n} \sum_{i=1}^n \frac{\hat{r}_i}{\pi(a_i, \mathbf{x}_i)} T\left(a_i \left(\sum_{j=1}^n v_j K(\mathbf{x}_i, \mathbf{x}_j) + b\right)\right). \quad (15)$$

Following a similar derivation to that used in the previous section, the convex subproblem at the $(t+1)$ 'th iteration of the d.c. algorithm is as follows,

$$\min_{v,b} \quad \frac{\lambda}{2} \sum_{i,j=1}^n v_i v_j K(\mathbf{x}_i, \mathbf{x}_j) + \frac{1}{n} \sum_{i=1}^n [\phi_1(u_i) \mathbb{I}(\hat{r}_i > 0) + \phi_0(u_i) \mathbb{I}(\hat{r}_i < 0)] \frac{|\hat{r}_i|}{\pi(a_i, \mathbf{x}_i)} + \sum_{i=1}^n \beta_i^{(t)} u_i,$$

where $u_i = a_i \left(\sum_{j=1}^n v_j K(\mathbf{x}_i, \mathbf{x}_j) + b \right)$. After solving the subproblem by L-BFGS, we update $\beta^{(t+1)}$ by (12). The procedure is repeated until β converges. When we obtain the solution $(\hat{v}, \hat{b})$, the decision function is $\hat{f}(\mathbf{x}) = \sum_{j=1}^n \hat{v}_j K(\mathbf{x}, \mathbf{x}_j) + \hat{b}$. Note that if we choose a linear kernel $K(\mathbf{x}, \mathbf{z}) = \mathbf{x}^T \mathbf{z}$, the obtained rule reduces to the previous linear rule. The most widely used nonlinear kernel in practice is the Gaussian Radial Basis Function (RBF) kernel, that is,

$$K_\sigma(\mathbf{x}, \mathbf{z}) = \exp\left(-\sigma^2 \|\mathbf{x} - \mathbf{z}\|^2\right),$$

where $\sigma > 0$ is a free parameter whose inverse $1/\sigma$ is called the width of K_σ .

2.4 A general framework for Residual Weighted Learning

We have proposed a method to identify the optimal ITR for continuous outcomes. However, in clinical practice, the endpoint outcome could also be binary, count, or rate. In this section, we provide a general framework to deal with all these types of outcomes.

The procedure is described as follows. First, estimate the conditional expected outcomes

given clinical covariates $\mathbf{X}$, $\hat{\mathbb{E}}\left(\frac{R}{2\pi(A, \mathbf{X})} | \mathbf{X}\right)$, by an appropriate regression model. It is equivalent to fitting a weighted regression model, in which each subject is weighted by

$\frac{1}{2\pi(A, \mathbf{X})}$. Second, we calculate the estimated residual $\hat{r}_i$ by comparing the observed outcome and expected outcome estimated in the first step. If the observed outcome is better than the expected one, the estimated residual $\hat{r}_i$ is positive, otherwise it is negative.

Specifically, if larger values of R are more desirable, as with our assumption for continuous

outcomes, $\hat{r}_i = r_i - \hat{\mathbb{E}}\left(\frac{R}{2\pi(A, \mathbf{X})} | \mathbf{X} = \mathbf{x}_i\right)$; if smaller values of R are preferred, e.g. the

number of adverse events, then $\hat{r}_i = \hat{\mathbb{E}}\left(\frac{R}{2\pi(A, \mathbf{X})} | \mathbf{X} = \mathbf{x}_i\right) - r_i$. Third, identify the decision function $\hat{f}(\mathbf{x})$ by using the estimated $\hat{r}_i$ in RWL.

The underlying idea is simple. Generally, the outcome R is a random variable depending on $\mathbf{X}$ and A . We estimate common effects of $\mathbf{X}$ by ignoring the treatment assignment A , and then all information on the heterogeneous treatment effects is contained in the residuals. As demonstrated previously, the benefits from using residuals include stabilizing the variance and controlling the treatment matching factor, both of which have a positive impact on finite sample performance. In previous sections, we discussed RWL for continuous outcomes in detail. We will provide two additional examples to illustrate the general framework.

The first example is for binary outcomes, *i.e.* $R \in \{0, 1\}$. Assume that $R = 1$ is desirable. We may fit a weighted main effects logistic regression model,

$$\mathbb{E}(R|\mathbf{X}) = \frac{\exp(\beta_0 + \mathbf{X}^T \beta)}{1 + \exp(\beta_0 + \mathbf{X}^T \beta)}.$$

The estimates $\hat{\beta}_0$ and $\hat{\beta}$ can be obtained numerically by statistical software, for example, the

glm function in R. The residual $\hat{r}_i$ is $r_i - \frac{\exp(\hat{\beta}_0 + \mathbf{X}_i^T \hat{\beta})}{1 + \exp(\hat{\beta}_0 + \mathbf{X}_i^T \hat{\beta})}$. Then we estimate the optimal ITR by RWL.

The second is for count/rate outcomes. We use the pulmonary exacerbation (PE) outcome of the cystic fibrosis data in Section 6 as an example. The outcome is the number of PEs during the study, and it is a rate variable. The observed data are $D_n = \{\mathbf{x}_i, a_i, r_i, t_i\}_{i=1}^n$, where r_i is the number of PEs in the duration t_i . We treat the count r_i as the outcome, and $(\mathbf{x}_i, t_i)$ as clinical covariates. We may fit a weighted main effects Poisson regression model,

$$\log(\mathbb{E}(R|\mathbf{X}, T)) = \beta_0 + \mathbf{X}^T \beta + \log(T).$$

Since fewer PEs are desirable, we compute $\hat{r}_i$ as $\exp(\hat{\beta}_0 + \mathbf{x}_i^T \hat{\beta} + \log(t_i)) - r_i$, which is the opposite of the raw residual in generalized linear models.

3 Theoretical Properties

In this section, we establish theoretical results for RWL. To the end, we assume that a sample $D_n = \{\mathbf{x}_i, a_i, r_i\}_{i=1}^n$ is independently drawn from a probability measure P on $\mathcal{X} \times \mathcal{A} \times \mathcal{M}$, where $\mathcal{X} \subset \mathbb{R}^p$ is compact, and $\mathcal{M} = [-M, M] \subset \mathbb{R}$. For any ITR $d: \mathcal{X} \rightarrow \mathcal{A}$, the risk is defined as

$$\mathcal{R}(d) = \mathbb{E} \left[\frac{R}{\pi(A, \mathbf{X})} \mathbb{I}(A \neq d(\mathbf{X})) \right].$$

The ITR that minimizes the risk is the Bayes rule $d^* = \arg \min_d \mathcal{R}(d)$, and the corresponding risk $\mathcal{R}^* = \mathcal{R}(d^*)$ is the Bayes risk. Recall that the Bayes rule is $d^*(\mathbf{x}) = \text{sign}(\mathbb{E}(R|\mathbf{X}=\mathbf{x}, A=1) - \mathbb{E}(R|\mathbf{X}=\mathbf{x}, A=-1))$.

In RWL, we substitute the 0-1 loss $\mathbb{I}(A \neq \text{sign}(f(\mathbf{X})))$ by the smoothed ramp loss, $T(Af(\mathbf{X}))$. Accordingly, we define the T -risk as

$$\mathcal{R}_{T,g}(f) = \mathbb{E} \left[\frac{R - g(\mathbf{X})}{\pi(A, \mathbf{X})} T(Af(\mathbf{X})) \right],$$

and, similarly, the minimal T -risk as $\mathcal{R}_{T,g}^* = \inf_f \mathcal{R}_{T,g}(f)$ and $f_{T,g}^* = \arg \min_f \mathcal{R}_{T,g}(f)$. In the theoretical analysis, we do not require g to be a regression fit of R . g can be any arbitrary measurable function. We may further assume that $(R - g(X))/\pi(A, X)$ is bounded almost surely.

The performance of an ITR associated with a real-valued function f is measured by the excess risk $\mathcal{R}(\text{sign}(f)) - \mathcal{R}^*$. In terms of the value function, we note that the excess risk is just $\mathcal{V}(d^*) - \mathcal{V}(\text{sign}(f))$. Similarly, we define the excess T -risk as $\mathcal{R}_{T,g}(f) - \mathcal{R}_{T,g}^*$.

Let $f_{D_n, \lambda_n} \in \mathcal{H}_K + \{1\}$, i.e. $f_{D_n, \lambda_n} = h_{D_n, \lambda_n} + b_{D_n, \lambda_n}$ where $h_{D_n, \lambda_n} \in \mathcal{H}_K$ and $b_{D_n, \lambda_n} \in \mathbb{R}$, be a global minimizer of the following optimization problem:

$$\min_{f=h+b \in \mathcal{H}_K + \{1\}} \frac{\lambda_n}{2} \|h\|_K^2 + \frac{1}{n} \sum_{i=1}^n \frac{r_i - g(x_i)}{\pi(a_i, x_i)} T(a_i f(x_i)).$$

Here we suppress g from the notation of f_{D_n, λ_n} , h_{D_n, λ_n} and b_{D_n, λ_n} .

The purpose of the theoretical analysis is to estimate the excess risk

$\mathcal{R}(\text{sign}(f_{D_n, \lambda_n})) - \mathcal{R}^*$ as the sample size n tends to infinity. Convergence rates will be derived under the choice of the parameter λ_n and conditions on the distribution P .

3.1 Fisher Consistency

We establish Fisher consistency of the decision function based on the smoothed ramp loss. Specifically, the following result holds:

Theorem 3.1. *For any measurable function f , if $f_{T,g}^*$ minimizes T -risk $\mathcal{R}_{T,g}(f)$, then the*

Bayes rule $d^(x) = \text{sign}(f_{T,g}^*(x))$ for all x such that*

$\mathbb{E}(R|X=x, A=1) \neq \mathbb{E}(R|X=x, A=-1)$. Furthermore, $\mathcal{R}_{T,g}(d^) = \mathcal{R}_{T,g}^*$.*

The proof is provided in the supplementary material. This theorem shows the validity of using the smoothed ramp loss as a surrogate loss in RWL.

3.2 Excess Risk

The following result establishes the relationship between the excess risk and excess T -risk. The proof can be found in the supplementary material.

Theorem 3.2. *For any measurable $f: \mathcal{X} \rightarrow \mathbb{R}$ and any probability distribution for (X, A, R) ,*

$$\mathcal{R}(\text{sign}(f)) - \mathcal{R}^* \leq \mathcal{R}_{T,g}(f) - \mathcal{R}_{T,g}^*.$$

This theorem shows that for any decision function f , the excess risk of f under the 0-1 loss is no larger than the excess risk of f under the smoothed ramp loss. It implies that if we obtain f by approximately minimizing $\mathcal{R}_{T,g}(f)$ so that the excess T -risk is small, then the risk of f is

close to the Bayes risk. In the next two sections, we investigate properties of f_{D_n, λ_n} through the excess T -risk.

3.3 Universal Consistency

In the literature of machine learning, a classification rule is called universally consistent if its expected error probability converges in probability to the Bayes error probability for all distributions underlying the data (Devroye et al. 1996). We extend this concept to ITRs. Precisely, given a sequence D_n of training data, an ITR d_n is said to be *universally consistent* if

$$\lim_{n \rightarrow \infty} \mathcal{R}(d_n) = \mathcal{R}^*$$

holds in probability for all probability measures P on $\mathcal{X} \times \mathcal{A} \times \mathcal{M}$. In this section, we will establish universal consistency of the rule $d_{D_n, \lambda_n} = \text{sign}(f_{D_n, \lambda_n})$. Before proceeding to the main result, we introduce some preliminary background on RKHSs, which is essential for our work.

Let K be a kernel, and $\mathcal{H}_K$ be its associated RKHS. We often use the quantity

$$C_K = \sup_{x \in \mathcal{X}} \sqrt{K(x, x)}.$$

In this article, we assume C_K is finite. For the Gaussian RBF kernel, $C_K = 1$. The assumption also holds for the linear kernel when $\mathcal{X} \subset \mathbb{R}^p$ is compact. By the reproducing property, it is easy to see that $\|f\|_\infty \leq C_K \|f\|_K$ for any $f \in \mathcal{H}_K$. A continuous kernel K on a compact metric space $\mathcal{X}$ is called *universal* if its associated RKHS $\mathcal{H}_K$ is dense in $C(\mathcal{X})$, i.e., for every function $g \in C(\mathcal{X})$ and all $\epsilon > 0$ there exists an $f \in \mathcal{H}_K$ such that $\|f - g\|_\infty < \epsilon$ (Steinwart and Christmann 2008, Definition 4.52). $C(\mathcal{X})$ is the space of all continuous functions $f: \mathcal{X} \rightarrow \mathbb{R}$ on the compact metric space $\mathcal{X}$ endowed with the usual supremum norm. The widely used Gaussian RBF kernel is an example of a universal kernel.

We are ready to present the main result of this section. The following theorem shows the convergence of the T -risk on the sample dependent function f_{D_n, λ_n} . We apply empirical process techniques to show consistency. The proof is provided in the supplementary material.

Theorem 3.3. Assume that we choose a sequence $\lambda_n > 0$ such that $\lambda_n \rightarrow 0$ and $n\lambda_n \rightarrow \infty$.

For a measurable function g , assume that $\frac{|R - g(\mathbf{X})|}{\pi(A, \mathbf{X})} \leq M_0$ almost surely. Then for any distribution P for $(\mathbf{X}, A, R)$, we have that in probability,

$$\lim_{n \rightarrow \infty} \mathcal{R}_{T, g}(f_{D_n}, \lambda_n) = \inf_{f \in \mathcal{H}_K + \{1\}} \mathcal{R}_{T, g}(f).$$

By Theorem 3.1, the Bayes rule d^* minimizes the T -risk $\mathcal{R}_{T,g}(f)$. Universal consistency follows if $\inf_{f \in \mathcal{H}_K + \{1\}} \mathcal{R}_{T,g}(f) = \mathcal{R}_{T,g}^*$ by Theorem 3.2. That is, d^* needs to be approximated by functions in the space $\mathcal{H}_K + \{1\}$. Clearly, this is not always true. A counterexample is when K is linear.

Let μ be the marginal distribution of X . Clearly, the rule d^* is measurable with respect to μ . Recall that a probability measure is regular if it is defined on the Borel sets. By Lusin's theorem, a measurable function can be approximated by a continuous function when μ is regular (Rudin 1987, Theorem 2.24). Thus the RKHS $\mathcal{H}_K$ of a universal kernel K is rich enough to provide arbitrarily accurate decision rules for all distributions. The finding is summarized in the following universal approximation lemma. The rigorous proof is provided in the supplementary material.

Lemma 3.4. *Let K be a universal kernel, and $\mathcal{H}_K$ be the associated RKHS. For all distributions P with regular marginal distribution μ on $\mathbf{X}$, we have*

$$\inf_{f \in \mathcal{H}_K + \{1\}} \mathcal{R}_{T,g}(f) = \mathcal{R}_{T,g}^*.$$

This immediately leads to universal consistency of RWL with a universal kernel.

Proposition 3.5. *Let K be a universal kernel, and $\mathcal{H}_K$ be the associated RKHS. Assume that we choose a sequence $\lambda_n > 0$ such that $\lambda_n \rightarrow 0$ and $n\lambda_n \rightarrow \infty$. For a measurable function*

g , assume that $\frac{|R - g(X)|}{\pi(A, X)} \leq M_0$ almost surely. Then for any distribution P for (X, A, R) with regular marginal distribution on $\mathbf{X}$, we have that in probability,

$$\lim_{n \rightarrow \infty} \mathcal{R}(\text{sign}(f_{D_n, \lambda_n})) = \mathcal{R}^*.$$

3.4 Convergence Rate

A consistent rule guarantees that by increasing the amount of data the rule can eventually learn the optimal decision with high probability. The next natural question is whether there are rules d_n with $\mathcal{R}(d_n)$ tending to the Bayes risk $\mathcal{R}^*$ at a specified rate for all distributions. Unfortunately, this is impossible as shown in the following theorem. The proof follows the arguments in Devroye et al. (1996, Theorem 7.2), which states in the classification problem the rate of convergence of any particular rule to the Bayes risk may be arbitrarily slow. We provide the outline of the proof in the supplementary material.

Theorem 3.6. *Assume there are infinite distinct points in $\mathcal{X}$. Let $\{c_n\}$ be a sequence of*

positive numbers converging to zero with $\frac{1}{16} \geq c_1 \geq c_2 \geq \dots$. For every sequence d_n of ITRs, there exists a distribution of (X, A, R) on $\mathcal{X} \times \mathcal{A} \times \mathcal{M}$ such that

$$\mathcal{R}(d_n) - \mathcal{R}^* \geq 2Mc_n.$$

This is a generic result, and applies to any algorithm for ITRs. It implies that rate of convergence studies for particular rules must necessarily be accompanied by conditions on (X, A, R) . Before moving back to RWL, we introduce the following quantity:

$$\mathcal{A}(\lambda) = \inf_{h \in \mathcal{H}_K, b \in \mathbb{R}} \frac{\lambda}{2} \|h\|_K^2 + \mathcal{R}_{T,g}(h+b) - \mathcal{R}_{T,g}^*.$$

The term $\mathcal{A}(\lambda)$ describes how well the infinite-sample regularized T -risk

$\frac{\lambda}{2} \|h\|_K^2 + \mathcal{R}_{T,g}(h+b)$ approximates the optimal T -risk $\mathcal{R}_{T,g}^*$. A similar quantity is called the *approximation error function* in the literature of learning theory (Steinwart and Christmann 2008, Definition 5.14). Since $\mathcal{A}(\lambda)$ is an infimum of affine linear functions, $\mathcal{A}(\lambda)$ is continuous. Thus

$$\lim_{\lambda \rightarrow 0} \mathcal{A}(\lambda) = \inf_{h \in \mathcal{H}_K, b \in \mathbb{R}} \mathcal{R}_{T,g}(h+b) - \mathcal{R}_{T,g}^*.$$

By Lemma 3.4, for a universal kernel K , $\lim_{\lambda \rightarrow 0} \mathcal{A}(\lambda) = 0$.

In order to establish the convergence rate, let us additionally assume that there exist constants $c > 0$ and $\beta \in (0, 1]$ such that

$$\mathcal{A}(\lambda) \leq c\lambda^\beta, \quad \forall \lambda > 0. \quad (16)$$

The assumption is standard in the literature of learning theory (Steinwart and Christmann 2008, p. 229). The following theorem is the main result in this section. This theorem is proved in the supplementary material using techniques of concentration inequalities developed in Bartlett and Mendelson (2002).

Theorem 3.7. *For any distribution P for (X, A, R) that satisfies the approximation error assumption (16), take $\lambda_n = n^{-\frac{1}{2\beta+1}}$. For a measurable function g , assume that*

$$\frac{|R - g(X)|}{\pi(A, X)} \leq M_0 \text{ almost surely. Then with probability at least } 1 - \delta,$$

$$\mathcal{R}(\text{sign}(f_{D_n, \lambda_n})) - \mathcal{R}^* \leq \tilde{c} \sqrt{\log(4/\delta)} n^{-\frac{\beta}{2\beta+1}},$$

where the constant $\tilde{c}$ is independent of δ and n , and $\tilde{c}$ decreases as M_0 decreases.

The resulting rate depends on the polynomial decay rate β of the approximation error, and the optimal rate is about $n^{-1/3}$ when β approaches 1. Theorem 3.7 also implies that even when the sample size n is fixed, an appropriately chosen g with a small bound M_0 would improve performance. It is easy to verify that g^* in (5) is the function g that minimizes

$\mathbb{E} \left(\frac{(R - g(\mathbf{X}))^2}{\pi(A, \mathbf{X})} \right)$. So when we choose g^* as in RWL, we may possibly obtain a smaller M_0 , as shown in Figure 1.

We are particularly interested in the Gaussian RBF kernel since it is widely used and is universal. It is important to understand the approximation error assumption (16) for the Gaussian RBF kernel. We first introduce the following “geometric noise” assumption for P on $(\mathbf{X}, A, R)$ (Steinwart and Scovel 2007). Let

$$\delta(\mathbf{x}) = \mathbb{E}(R|\mathbf{X}=\mathbf{x}, A=1) - \mathbb{E}(R|\mathbf{X}=\mathbf{x}, A=-1).$$

Define $\mathcal{X}^+ = \{\mathbf{x} \in \mathcal{X} : \delta(\mathbf{x}) > 0\}$, $\mathcal{X}^- = \{\mathbf{x} \in \mathcal{X} : \delta(\mathbf{x}) < 0\}$, and $\mathcal{X}^0 = \{\mathbf{x} \in \mathcal{X} : \delta(\mathbf{x}) = 0\}$. Now we define a distance function $\mathbf{x} \mapsto \tau_{\mathbf{x}}$ by

$$\tau_{\mathbf{x}} = \begin{cases} \tilde{d}(\mathbf{x}, \chi^0 \cup \chi^+) & \text{if } \mathbf{x} \in \mathcal{X}^-, \\ \tilde{d}(\mathbf{x}, \chi^0 \cup \chi^-) & \text{if } \mathbf{x} \in \mathcal{X}^+, \\ 0 & \text{if } \mathbf{x} \in \mathcal{X}^0, \end{cases}$$

where $\tilde{d}(\mathbf{x}, A)$ denotes the distance of $\mathbf{x}$ to a set A with respect to the Euclidean norm. Now we present the geometric noise condition for distributions.

Definition 3.8. Let $\mathcal{X} \subset \mathbb{R}^p$ be compact and P be a probability measure of $(\mathbf{X}, A, R)$ on $\mathcal{X} \times \mathcal{A} \times \mathcal{M}$. We say that P has geometric noise exponent $q > 0$ if there exists a constant $C > 0$ such that

$$\int_{\mathcal{X}} \exp\left(-\frac{\tau_{\mathbf{x}}^2}{t}\right) |\delta(\mathbf{x})| \mu(d\mathbf{x}) \leq Ct^{qp/2}, \quad t > 0, \quad (17)$$

where μ is the marginal measure on $\mathbf{X}$.

The integral condition (17) describes the concentration of the measure $|\delta(\mathbf{x})|d\mu$ near the decision boundary in the sense that the less the measure is concentrated in this region the larger the geometric noise exponent can be chosen. The following lemma shows that the geometric noise condition can be used to guarantee the approximation error assumption (16) when the parameter σ of the Gaussian RBF kernel K_{σ} is appropriately chosen.

Lemma 3.9. Let $\mathcal{X}$ be the closed unit ball of the Euclidean space $\mathbb{R}^p$, and P be a distribution on $\mathcal{X} \times \mathcal{A} \times \mathcal{M}$ that has geometric noise exponent $0 < q < \infty$ with constant C in (17). Let $\sigma = \lambda^{-\frac{1}{(q+1)p}}$. Then for the Gaussian RBF kernel K_{σ} , there is a constant $c > 0$ depending only on the dimension p , the geometric noise exponent q and constant C , such that for all $\lambda > 0$ we have

$$\mathcal{A}(\lambda) \leq c\lambda^{q/(q+1)}.$$

The proof provided in the supplementary material follows the idea from Theorem 2.7 in Steinwart and Scovel (2007) by replacing their $2\eta(\mathbf{x}) - 1$ with $\delta(\mathbf{x})$. The key point is the following inequality, for all measurable f ,

$$\mathcal{R}_{T,g}(f) - \mathcal{R}_{T,g}^* \leq 2\mathbb{E}(|\delta(\mathbf{x})| \cdot |f - d^*|),$$

which is the counterpart of Equation (24) in Steinwart and Scovel (2007).

Now we are ready to present the convergence rate of RWL with Gaussian RBF kernel by combining Theorem 3.7 and Lemma 3.9:

Proposition 3.10. *Let $\mathcal{X}$ be the closed unit ball of the Euclidean space $\mathbb{R}^p$, and P be a distribution on $\mathcal{X} \times \mathcal{A} \times \mathcal{M}$ that has geometric noise exponent $0 < q < \infty$ with constant C*

in (17). For a measurable function g , assume that $\frac{|R - g(\mathbf{X})|}{\pi(\mathbf{A}, \mathbf{X})} \leq M_0$ almost surely.

Consider a Gaussian RBF kernel K_{σ_n} . Take $\lambda_n = n^{-\frac{q}{3q+1}}$ and $\sigma_n = \lambda_n^{-\frac{1}{(q+1)p}}$. Then with probability at least $1 - \delta$,

$$\mathcal{R}(\text{sign}(f_{D_n, \lambda_n})) - \mathcal{R}^* \leq \tilde{c} \sqrt{\log(4/\delta)n^{-\frac{q}{3q+1}}},$$

where the constant $\tilde{c}$ is independent of δ and n , and $\tilde{c}$ decreases as M_0 decreases.

The optimal rate for the risk is about $n^{-1/3}$ when the geometric noise exponent q is sufficiently large. A better rate may be achieved by using techniques in Steinwart and Scovel (2007), but it is beyond the scope of this work.

4 Variable selection for RWL

Variable selection is an important area in modern statistical research. Gunter et al. (2011) distinguished prescriptive covariates from predictive covariates. The latter refers to covariates which are related to the prediction of outcomes, and the former refers to covariates which help to prescribe optimal ITRs. Clearly, prescriptive covariates are predictive too. In practice, a large number of clinical covariates are often available for estimating optimal ITRs. However, many of them might not be prescriptive. Thus careful variable selection could lead to better performance.

4.1 Variable selection for linear RWL

There is a vast body of literature on variable selection for linear classification in statistical learning. For instance, Zhu et al. (2003) and Fung and Mangasarian (2004) investigated the support vector machine (SVM) using the ℓ_1 -norm penalty (Tibshirani 1994). Wang et al. (2008) applied the elastic-net penalty (Zou and Hastie 2005) for variable selection with an SVM-like method. In this article, we replace the ℓ_2 -norm penalty in RWL with the elastic-net penalty,

$$\lambda_1 \|w\|_1 + \frac{\lambda_2}{2} w^T w,$$

where $\|w\|_1 = \sum_{j=1}^p |w_j|$ is the ℓ_1 -norm. As a hybrid of the ℓ_1 -norm and ℓ_2 -norm penalties, the elastic-net penalty retains the variable selection features of the ℓ_1 -norm penalty, and tends to provide similar estimated coefficients for highly correlated variables, *i.e.* the grouping effect, as the ℓ_2 -norm penalty does. Hence, highly correlated variables are selected or removed together.

The elastic-net penalized linear RWL aims to minimize

$$\lambda_1 \|w\|_1 + \frac{\lambda_2}{2} w^T w + \frac{1}{n} \sum_{i=1}^n \frac{\hat{r}_i}{\pi(a_i, x_i)} T(a_i (w^T x_i + b)),$$

where $\lambda_1 (> 0)$ and $\lambda_2 (> 0)$ are regularization parameters. We still apply the d.c. algorithm to solve the above optimization problem. Following a similar decomposition, the convex subproblem at the $(t+1)$ 'th iteration of the d.c. algorithm is as follows,

$$\min_{w, b} \lambda_1 \|w\|_1 + \frac{\lambda_2}{2} w^T w + \frac{1}{n} \sum_{i=1}^n [\phi_1(u_i) \mathbb{I}(\hat{r}_i > 0) + \phi_0(u_i) \mathbb{I}(\hat{r}_i < 0)] \frac{|\hat{r}_i|}{\pi(a_i, x_i)} + \sum_{i=1}^n \beta_i^{(t)} u_i,$$

(18)

where $u_i = a_i(w^T x_i + b)$, and $\beta_i^{(t)}$ is computed by (12). The objective function is a sum of a smooth function and a non-smooth ℓ_1 -norm penalty. Since L-BFGS can only deal with smooth objective functions, it may not work here. Instead we use projected scaled subgradient (PSS) algorithms (Schmidt 2010), which are extensions of L-BFGS to the case of optimizing a smooth function with an ℓ_1 -norm penalty. Several PSS algorithms are proposed in Schmidt (2010). Among them, the Gafni-Bertsekas variant is particularly effective for our purpose. After solving the subproblem, we update $\beta_i^{(t+1)}$ by (12). The procedure is repeated until β converges. The obtained decision function is $\hat{f}(x) = \hat{w}^T x + \hat{b}$, and thus the estimated optimal ITR is the sign of $\hat{f}(x)$.

4.2 Variable selection for RWL with nonlinear kernels

Many researchers have noticed that nonlinear kernel methods may perform poorly when there are irrelevant covariates presented in the data (Weston et al. 2000; Grandvalet and Canu 2002; Lin and Zhang 2006; Lafferty and Wasserman 2008). For optimal treatment regimes, we may suffer a similar problem when using a nonlinear kernel. For the Gaussian RBF kernel, we assume that the geometric noise condition in Definition 3.8 is satisfied with exponent q . When several additional non-descriptive covariates are included in the data, the geometric noise exponent is decreased according to (17), and hence so is the convergence

rate in Proposition 3.10. The presence of non-descriptive covariates thus can deteriorate the performance of RWL.

The variable selection approach proposed in this section is inspired by the KNIFE algorithm (Allen 2013), where a set of scaling factors on the covariates are employed to form a regularized loss function. The idea of scaling covariates within the kernel has appeared in several methods from the machine learning community (Weston et al. 2000; Grandvalet and Canu 2002; Rakotomamonjy 2003; Argyriou et al. 2006).

We take the Gaussian RBF kernel as an example. Define the covariates-scaled Gaussian RBF kernel,

$$K_{\eta}(\mathbf{x}, \mathbf{z}) = \exp \left(- \sum_{j=1}^p \eta_j (x_j - z_j)^2 \right),$$

where $\boldsymbol{\eta} = (\eta_1, \dots, \eta_p)^T \geq \mathbf{0}$. Here each covariate x_j is scaled by $\sqrt{\eta_j}$. Setting $\eta_j = 0$ is equivalent to discarding the j th covariate. The hyperparameter σ in the original Gaussian RBF kernel is absorbed to the scaling factors. Similar with KNIFE, we seek $(\hat{v}, \hat{b}, \hat{\eta})$ to minimize the following optimization problem:

$$\begin{aligned} \min_{\mathbf{v}, \mathbf{b}, \boldsymbol{\eta}} \quad & \lambda_1 \|\boldsymbol{\eta}\|_1 + \frac{\lambda_2}{2} \sum_{i,j=1}^n v_i v_j K_{\eta}(\mathbf{x}_i, \mathbf{x}_j) + \frac{1}{n} \sum_{i=1}^n \frac{\hat{r}_i}{\pi(a_i, \mathbf{x}_i)} T \left(a_i \left(\sum_{j=1}^n v_j K_{\eta}(\mathbf{x}_i, \mathbf{x}_j) + b \right) \right), \\ \text{subject to} \quad & \boldsymbol{\eta} \geq \mathbf{0}, \end{aligned} \quad (19)$$

where $\lambda_1 (> 0)$ and $\lambda_2 (> 0)$ are regularization parameters. Compared with previous nonlinear RWL optimization problem in (15), (19) has an ℓ_1 -norm penalty on scaling factors. The objective function is singular at $\boldsymbol{\eta} = \mathbf{0}$ due to nonnegativity of scaling factors. So (19) may produce zero solutions for some of the $\boldsymbol{\eta}$, and hence performs variable selection.

We adopt a similar decomposition trick that has been used several times before. The subproblem at the $(t+1)$ 'th iteration is

$$\begin{aligned} \min_{\mathbf{v}, \mathbf{b}, \boldsymbol{\eta}} \quad & \lambda_1 \|\boldsymbol{\eta}\|_1 + \frac{\lambda_2}{2} \sum_{i,j=1}^n v_i v_j K_{\eta}(\mathbf{x}_i, \mathbf{x}_j) + \frac{1}{n} \sum_{i=1}^n [\phi_1(u_i) \mathbb{I}(\hat{r}_i > 0) + \phi_0(u_i) \mathbb{I}(\hat{r}_i < 0)] \frac{|\hat{r}_i|}{\pi(a_i, \mathbf{x}_i)} + \sum_{i=1}^n \beta_i^{(t)} u_i, \\ \text{subject to} \quad & \boldsymbol{\eta} \geq \mathbf{0}, \end{aligned} \quad (20)$$

where $u_i = a_i \left(\sum_{j=1}^n v_j K_{\eta}(\mathbf{x}_i, \mathbf{x}_j) + b \right)$, and $\beta_i^{(t)}$ is computed by (12). Recall that the d.c. algorithm is a special case of the Majorize-Minimization (MM) algorithm. Although with the covariate-scaled kernel the subproblem is nonconvex, we still apply a similar iterative procedure as in the d.c. algorithm, *i.e.* solving a sequence of subproblems (20), but based on the MM algorithm. The subproblem (20) is a smooth optimization problem with box constraints. In this article, we use L-BFGS-B (Byrd et al. 1995; Morales and Nocedal 2011), which extends L-BFGS to handle simple box constraints on variables. After solving the subproblem (20), we update u_j and β . The procedure is repeated until convergence. Then the obtained decision function is $\hat{f}(\mathbf{x}) = \sum_{i=1}^n \hat{v}_i K_{\hat{\eta}}(\mathbf{x}_i, \mathbf{x}) + \hat{b}$, and thus the estimated optimal ITR is the sign of $\hat{f}(\mathbf{x})$. The covariates with nonzero scaling factors $\hat{\eta}$ are identified to be important in estimating ITRs.

This variable selection technique can be applied to other nonlinear kernels, such as the polynomial kernel, and even to the linear kernel. However for the linear kernel, we recommend the approach in the prior section, where the subproblem (18) is a convex optimization problem with $p+1$ variables. Note that the subproblem (20) in this section is nonconvex even for the linear kernel, and there are $n + p + 1$ variables for the optimization problem. This is much more challenging.

5 Simulation studies

We carried out extensive simulation studies to investigate finite sample performance of the proposed RWL methods.

We first evaluated the performance of RWL methods with low-dimensional covariates. In the simulations, we generated 5-dimensional vectors of clinical covariates $x_1, \dots, x_5$, consisting of independent uniform random variables $U(-1, 1)$. The treatment A was generated from $\{-1, 1\}$ independently of X with $P(A = 1) = 0.5$. That is, $\pi(a, \mathbf{x}) = 0.5$ for all a and $\mathbf{x}$. The response R was normally distributed with mean $Q_0 = \mu_0(\mathbf{x}) + \delta_0(\mathbf{x}) \cdot a$ and standard deviation 1, where $\mu_0(\mathbf{x})$ is the common effect for clinical covariates, and $\delta_0(\mathbf{x}) \cdot a$ is the interaction between treatment and clinical covariates. We considered five scenarios with different choices of $\mu_0(\mathbf{x})$ and $\delta_0(\mathbf{x})$:

- (0) $\mu_0(\mathbf{x}) = 1 + x_1 + x_2 + 2x_3 + 0.5x_4; \delta_0(\mathbf{x}) = 0.4(x_2 - 0.25x_1^2 - 1).$
- (1) $\mu_0(\mathbf{x}) = 1 + x_1 + x_2 + 2x_3 + 0.5x_4; \delta_0(\mathbf{x}) = 1.8(0.3 - x_1 - x_2).$
- (2) $\mu_0(\mathbf{x}) = 1 + x_1 + x_2 + 2x_3 + 0.5x_4; \delta_0(\mathbf{x}) = 1.3(x_2 - 2x_1^2 + 0.3).$
- (3) $\mu_0(\mathbf{x}) = 1 + x_1^2 + x_2^2 + 2x_3^2 + 0.5x_4^2; \delta_0(\mathbf{x}) = 3.8(0.8 - x_1^2 - x_2^2).$
- (4) $\mu_0(\mathbf{x}) = 1 + x_1^2 + x_2^2 + 2x_3^2; \delta_0(\mathbf{x}) = 10(1 - x_1^2 - x_2^2)(x_1^2 + x_2^2 - 0.2).$

Scenario 0 is similar with the second scenario in Zhao et al. (2012). When x_2 is restricted to $[-1, 1]$, the treatment arm -1 is always better than the treatment arm 1 on average. Thus the true optimal ITR in Scenario 0 is to assign all subjects to arm -1 . The decision boundaries

for the remaining four scenarios are determined by x_1 and x_2 only. The scenarios have different decision boundaries in truth. The decision boundary is a line in Scenario 1, a parabola in Scenario 2, a circle in Scenario 3, and a ring in Scenario 4. Their true optimal ITRs are illustrated in Figure 1 in the supplementary material. It is unclear how often a circle or ring boundary structure will occur in practice, although general nonlinear boundaries are likely to arise frequently. The purpose of including these in the simulations is to verify that the proposed approach can handle even the most difficult boundary structures. The coefficients in Scenarios 0 ~ 3 were chosen to reflect a medium effect size according to Cohen's d index (Cohen 1988); and the coefficients in Scenario 4 reflect a small effect size. The Cohen's d index is defined as the standardized difference in mean responses between two treatment arms, that is,

$$es = \frac{|\mathbb{E}(R|A=1) - \mathbb{E}(R|A=-1)|}{\sqrt{[\text{Var}(R|A=1) + \text{Var}(R|A=-1)]/2}}.$$

We compared the performance of the following five methods:

- (1) The proposed RWL using the Gaussian RBF kernel (RWL-Gaussian). Residuals were estimated by a linear main effects model.
- (2) The proposed RWL using the linear kernel (RWL-Linear). Residuals were estimated by a linear main effects model.
- (3) OWL proposed in Zhao et al. (2012) using the Gaussian RBF kernel (OWL-Gaussian).
- (4) OWL using the linear kernel (OWL-Linear).
- (5) ℓ_1 -PLS proposed by Qian and Murphy (2011).

The OWL methods were reviewed in the method section. Since the OWL methods can only handle nonnegative outcomes, we subtracted $\min\{r_j\}$ from all outcome responses. This is essentially the approach used in Zhao et al. (2012) [personal communication]. In the simulation studies, ℓ_1 -PLS approximated $\mathbb{E}(R|X, A)$ using the basis function set $(1, X, A, XA)$, and applied LASSO (Tibshirani 1994) for variable selection. The optimal ITR was determined by the sign of the difference between the estimated $\mathbb{E}(R|X, A=1)$ and $\mathbb{E}(R|X, A=-1)$.

For each scenario, we considered two sample sizes for training datasets: $n = 100$ and $n = 400$, and repeated the simulation 500 times. All five methods had at least one tuning parameter. For example, there was one tuning parameter, λ , in RWL-linear; and two tuning parameters, λ and σ in RWL-Gaussian. In the simulations, we applied a 10-fold cross-validation procedure to tune parameters. For RWL-Linear, we searched over a pre-specified finite set of λ 's to select the best one maximizing the average of estimated values from validation data; while for RWL-Gaussian, we searched over a pre-specified finite set of (λ, σ) 's to select the best pair. The selection of tuning parameters in ℓ_1 -PLS and OWL followed similarly. In the comparisons, the performances of five methods were evaluated by two criteria: the first criterion is the value function under the estimated optimal ITR when we

applied to an independent and large test dataset; the second criterion is the misclassification error rate of the estimated optimal ITR from the true optimal ITR on the large test dataset. Specifically, a test set with 10,000 observations was simulated to assess the performance. The estimated value function under any ITR d is given by

$\mathbb{P}_n^* [R\mathbb{I}(A=d(\mathbf{X})) / \pi(A, \mathbf{X})] / \mathbb{P}_n^* [\mathbb{I}(A=d(\mathbf{X})) / \pi(A, \mathbf{X})]$ (Murphy 2005), where $\mathbb{P}_n^*$ denotes the empirical average using the test data; the misclassification rate under any ITR d is given by $\mathbb{P}_n^* [\mathbb{I}(d(\mathbf{X}) \neq d^*(\mathbf{X}))]$, where d^* is the true optimal ITR which is known when generating the simulated data.

The simulation results are presented in Table 1. We report the mean and standard deviation of value functions and misclassification rates over 500 replications. In Scenario 0, the true optimal ITR would assign all subjects to treatment arm -1. When the sample size was small ($n = 100$), ℓ_1 -PLS, OWL-Linear and RWL-Linear showed similar performance, while OWL-Gaussian and RWL-Gaussian were slightly worse. Note that when x_2 can go beyond 1, the heterogeneous treatment effect does exist with non-linear (parabola) decision boundary. The deteriorated performance of OWL-Gaussian and RWL-Gaussian may be due to unexpected extrapolation. When the sample size was large ($n = 400$), all methods performed similarly, and assigned almost all subjects to arm -1. Scenario 1 was a linear example. The model specification in ℓ_1 -PLS was correct, and this method performed very well. The proposed RWL methods, with either the linear kernel or the Gaussian RBF kernel, yielded similar performance with ℓ_1 -PLS, and were much better than OWL methods. Another advantage of using RWL versus OWL was also reflected by a smaller variance in value functions. The idea of using residuals instead of the original outcomes is able to stabilize the variance, as we discussed in the method section.

The conditional mean outcomes and decision boundaries in the remaining three scenarios were nonlinear. ℓ_1 -PLS, RWL-Linear, and OWL-Linear were misspecified, and hence they did not perform very well in these three scenarios. We focus on the comparison between RWL-Gaussian and OWL-Gaussian. These two methods equipped with the Gaussian RBF kernel should have the ability to capture the nonlinear structure of decision boundaries. In Scenarios 2 and 3, RWL-Gaussian outperformed OWL-Gaussian in terms of achieving higher values and smaller variances. It was challenging to find the optimal ITR in Scenario 4 since the decision boundary was complicated and the effective size was small. When the sample size was small ($n = 100$), the performance of RWL-Gaussian was only slightly better than that of OWL-Gaussian, but with larger variance. It might be too hard to learn the complicated decision boundary from the limited sample. However, when the sample size increased, RWL-Gaussian showed significantly better performance than other methods. The excellent performance of RWL-Gaussian confirms its power in finding the optimal ITR. Although using the Gaussian RBF kernel, the performance of OWL-Gaussian was not comparable with that of RWL-Gaussian. In Scenario 2, OWL-Gaussian performed even worse than ℓ_1 -PLS and RWL-Linear, both of which can only detect linear boundaries. While in Scenario 4, OWL-Gaussian gave similar performance as OWL-Linear. Zhao et al. (2012) also reported similar finding in their simulation studies that there are no large differences in the performance between OWL-Linear and OWL-Gaussian.

As we demonstrated in Section 2.2, OWL methods tend to favor the treatment assignments that subjects actually received, and this behavior deteriorates finite sample performance. We calculated treatment matching factors of these five methods on the training data. They are shown in Table 1 in the supplementary material. We expect the treatment matching factor to be close to 1. However, the treatment matching factors for OWL methods were greater than 1. OWL-Gaussian achieved the largest treatment matching factors, especially when the sample size was small. This may partially explain why OWL-Gaussian may not outperform OWL-Linear even when the decision boundary is nonlinear. Whereas for RWL methods, the treatment matching factors were close to 1. By appropriately controlling the treatment matching factors, the aim of RWL methods is to maximize the estimate in (8), and hence to improve finite sample performance.

We applied a linear main effects model to estimate residuals. The model was correctly specified for Scenarios 0, 1 and 2, and misspecified for Scenarios 3 and 4. However, the superior performance of RWL methods in Scenarios 3 and 4 demonstrates the robustness of our methods to residual estimates. We also considered a null model, which was misspecified for all scenarios, to estimate residuals. The results were only slightly worse than those shown in Table 1 for Scenarios 1 and 2 when $n = 100$, and other scenarios were similar especially when the sample size increased [results not shown].

We then evaluated the performance of RWL methods with moderate-dimensional clinical covariates. We adopted the same data generating procedure as in the above five scenarios except that the dimension of clinical covariates was increased from 5 to 50 by adding 45 independent uniform random variables $U(-1, 1)$. Thus among all these 50 clinical covariates, only x_1 and x_2 are attributed to the decision boundaries for Scenarios 1 ~ 4. We repeated the simulations, and the results are shown in Table 2. In Scenario 0, all methods performed similar, and assigned almost all subjects to the treatment arm -1 when the sample size was large ($n = 400$). In Scenario 1, ℓ_1 -PLS performed very well because of correct model specification and inside variable selection techniques. Although RWL methods showed better performance than OWL methods, they were not comparable to ℓ_1 -PLS. The decision boundaries in Scenarios 2 ~ 4 are nonlinear. Both RWL and OWL failed to detect the decision boundaries in these scenarios. They produced similar performance with either the linear kernel or Gaussian RBF kernel. RWL and OWL methods with Gaussian RBF kernel may perform particularly poorly when many non-descriptive covariates are present.

We tested the performance of variable selection approaches described in Section 4 for the linear kernel and Gaussian RBF kernel. They are called RWL-VS-Linear and RWL-VS-Gaussian, respectively, in this article. There were two tuning parameters, λ_1 and λ_2 , for both approaches. We applied a 10-fold cross-validation procedure to tune parameters from a pre-specified finite set of (λ_1, λ_2) 's. The simulation results are also shown in Table 2. In Scenario 0, both RWL-VS methods showed similar performance as RWL methods. In Scenario 1, both RWL-VS-Linear and RWL-VS-Gaussian improved the performance of RWL methods. When the sample size was 400, the performance was as good as that of ℓ_1 -PLS. For Scenarios 2 ~ 4, we only present results from RWL-VS-Gaussian. When the sample size was small ($n = 100$), the performance was slightly better (in Scenarios 2 and 3) than, or similar (in scenario 4) as that of RWL-Gaussian. It is well known that large sample

sizes are needed to find interactions in models. So for similar reasons, we may expect large samples to be necessary to find and confirm the existence of heterogeneity of treatment effects, and accurately define the decision boundary. In the simulations, when the sample size increased, the performance of RWL-VS-Gaussian improved significantly. Variable selection is necessary for RWL when there are many non-descriptive clinical covariates.

6 Data analysis

We applied the proposed methods to analyze data from the EPIC randomized clinical trial (Treggiari et al. 2011). The trial was designed to determine the optimal anti-pseudomonal treatment strategy in children with cystic fibrosis (CF) who recently acquired *Pseudomonas aeruginosa* (*Pa*). Newly identified *Pa* was defined as the first lifetime documented *Pa* positive culture within 6 months of baseline or a *Pa* positive culture within 6 months of baseline after a two-year absence of *Pa* culture-positivity. Other eligibility criteria have been previously reported (Treggiari et al. 2009). A total of 304 patients ages 1~12 were randomized in a 1:1 ratio to one of two maintenance treatment strategies: cycled therapy (treatment with anti-pseudomonal therapy provided in quarterly cycles regardless of *Pa* positivity) and culture-based therapy (treatment only in response to identification of *Pa* positive cultures from quarterly cultures). All patients regardless of randomization strategy received an initial course of anti-pseudomonal therapy in response to their *Pa* positive culture at eligibility and proceeded with treatment according to their randomization strategy. In this article, we considered two endpoints over the course of the 18-month study. One is the number of *Pa* positive cultures from scheduled follow-up quarterly cultures, and the other is the number of pulmonary exacerbations (PE) requiring any (intravenous, inhaled, or oral) antibiotic use or hospitalization during the study. In the original clinical trial, there was no significant difference between the two maintenance treatment strategies at the population level for either the *Pa* outcome (p-value 0.222) or PE outcome (p-value 0.280).

We considered 10 baseline clinical covariates, including age, gender, F508del genotype, weight, height, BMI, first documented lifetime *Pa* positive culture at eligibility versus positive after a two-year history of negative cultures (neverpapos), use of antibiotics within 6 months prior to baseline (anyabxhist), *Pa* positive at the baseline visit versus *Pa* positive within 6 months prior to baseline only (pabl), and positive versus negative baseline *S. aureus* culture status (sabl). There were three levels (homozygous, heterozygous and other) for F508del genotype. We coded genotype as two dichotomous covariates, homozygous and heterozygous, each referenced in relation to the “othe” genotype category. Four patients did not give consent to have their data put into the databank. We also excluded 17 patients with missing covariate values from the analyses. The data used here consisted of 283 patients: 141 in the cycled therapy, and 142 in the culture-based therapy. Each clinical covariate was linearly scaled to the range $[-1, +1]$, as recommended in Hsu et al. (2003).

For the *Pa*-related endpoint, we used the ratio of the number of *Pa* negative cultures to the number of *Pa* cultures over the follow up period as the outcome. We were seeking an ITR to reduce the number of *Pa* positive cultures, so larger outcomes were preferred. We compared seven methods: ℓ_1 -PLS, OWL-Linear, OWL-Gaussian, RWL-Linear, RWL-Gaussian, RWL-VS-Linear, and RWL-VS-Gaussian. ℓ_1 -PLS used the basis function set $(1, X, A, XA)$ in the

regression model. We treated the outcome as continuous, and used a linear main effects model to estimate residuals for RWL methods. We used a 10-fold cross-validation procedure to tune parameters. The estimated rule was then used to predict the optimal treatment for each patient. The predicted treatment allocation is shown in Table 3. To evaluate the performance of estimated rules, we again carried out a 10-fold cross validation procedure. The data were randomly partitioned into 10 roughly equal-sized parts. We estimated the ITR on nine parts of the data using the tuned parameters, and predicted optimal treatments on the part left out. We repeated the procedure nine more times to predict the other parts, and obtained the predicted treatment for each patient. The first evaluation criterion was the value function, which is given by $\mathbb{P}_n [R\mathbb{I}(A=Pred) / \pi(A, \mathbf{X})] / \mathbb{P}_n [\mathbb{I}(A=Pred) / \pi(A, \mathbf{X})]$, where $\mathbb{P}_n$ denotes the empirical average over a fold in cross validation and $Pred$ is the predicted treatment in the cross validation procedure. The second evaluation criterion was related to significance tests. The test was performed when a 10-fold cross validation procedure is finished, that is, when we had the predicted treatments for all patients. By comparing the predicted treatments with the treatments that were actually assigned to patients, we divided patients into two groups: those who followed the estimated optimal rule and those who did not follow the rule. We tested the difference between the two groups using the two-sample t-test to see whether the group that followed the estimated rule was better than the other group at the significance level $\alpha = 0.05$. The whole procedure was repeated 100 times with different fold partitions in the cross validation. We obtained 1000 value functions and 100 p-values. The mean and standard deviation of these value functions, the proportion of significant p-values and median of these p-values are also presented in Table 3. As reference rules, we also considered two fixed treatment rules: assign all patients to (1) cycled therapy arm or (2) culture-based therapy arm. We show in Table 3 their value functions from repeated 10-fold cross validation procedures and one-sided p-values from two-sample t-tests. Note that the p-values for fixed rules were not changed during the cross validation procedure. ℓ_1 -PLS showed the best performance, and identified that one covariate, baseline Pa status, is important in estimating ITRs. The significance tests also confirmed its superior performance. Both OWL methods failed to detect the heterogeneity of treatment effects. They assigned all patients to the cycled therapy arm. Although the performances of RWL methods (without variable selection) were only slightly better, on average, than those of OWL methods, they did detect some heterogeneity across treatments. Variable selection further improved the performance in terms of achieving higher values and better significance tests. Moreover, both RWL-VS-Linear and RWL-VS-Gaussian selected the same covariate as ℓ_1 -PLS.

Our results are consistent with results from the trial suggesting that the subgroup of patients who were Pa negative at the baseline visit (albeit positive within 6 months of baseline to meet eligibility) may have more greatly benefited from cycled therapy to suppress Pa positivity during the trial. However, it is important to note clinically that further analyses of the trial data demonstrated the comparable effectiveness using culture-based therapy as compared to cycled therapy in reducing the overall prevalence of Pa positive cultures at the end of the trial.

The number of PEs during the study was a rate variable. We were seeking an ITR to lower the number of PEs, so fewer PEs were desirable. We compared several methods: PoissonReg, ℓ_1 -PoissonReg, OWL-Linear, OWL-Gaussian, RWL-Linear, RWL-Gaussian, RWL-VS-Linear, and RWL-VS-Gaussian. PoissonReg fitted a Poisson regression model using the basis function set $(1, \mathbf{X}, \mathbf{A}, \mathbf{X} \mathbf{A})$, computed the predicted number of PEs for a particular patient at each arm, and assigned the treatment with smaller prediction to this patient. ℓ_1 -PoissonReg was similar with PoissonReg, but used an ℓ_1 penalized term to perform variable selection. For OWL methods, we used the opposite of individual annual PE rate as the continuous outcome. The annual PE rate for the i -th patient is defined as

$\frac{r_i}{t_i} * 365.25$, where r_i is the number of PEs for the i -th patient, and t_i is the duration (in days) when the i -th patient stayed in the study. For RWL methods, we used a main effects Poisson regression model to estimate residuals, as described in Section 2.4. A 10-fold cross-validation was used to tune parameters. The predicted treatment allocation is shown in Table 4. A similar evaluation procedure as that for the *Pa* endpoint was also performed, where we used instead the annual PE rate for the group that followed the rule as the first evaluation

criterion. The group-wise annual PE rate is given as $\frac{\mathbb{P}_n [R\mathbb{I}(A=Pred)]}{\mathbb{P}_n [T\mathbb{I}(A=Pred)]} * 365.25$, where R is the number of PEs, T is the duration (in days) in the study, $\mathbb{P}_n$ denotes the empirical average over a fold in cross validation and *Pred* is the predicted treatment assignment in the cross validation procedure. For the second criterion, a Poisson regression model was used to test whether the group that followed the estimated rule had fewer PEs than the other group that did not follow the estimated rule. The results are presented in Table 4. We also considered two fixed rules as references. All RWL methods outperformed OWL and Poisson regression methods in terms of achieving better PE rates and better significance test results. ℓ_1 -PoissonReg identified three important covariates, baseline *Pa* status, baseline *S. aureus* status and height, for treatment assignment.

For this case, RWL-VS-Linear and RWL-VS-Gaussian did not improve performances comparing with their counterparts without variable selection. RWL-VS-Linear did not discard any clinical covariates. As its tuned parameter λ_1 was very small, RWL-VS-Linear gave a similar performance as RWL-Linear. The estimated rule from RWL-Linear was determined by the following decision function F ,

$$\begin{aligned} F = & 1.43 * \mathbb{I}(gender=female) + 0.07 * age \\ & - 0.58 * \mathbb{I}(neverpapos \\ & =yes) - 1.01 * \mathbb{I}(pabl='+') \\ & - 0.17 * \mathbb{I}(sable='+') \\ & + 0.02 * weight \\ & + 0.01 * height \\ & - 0.17 * \mathbb{I}(anyabxhist \\ & =yes) + 0.11 * bmi \\ & - 0.90 * \mathbb{I}(genotype \\ & =homozygous) - 0.83 * \mathbb{I}(genotype=heterozygous) - 2.90 \end{aligned}$$

The corresponding ITR is that if $F > 0$, assign cycled therapy to this patient, otherwise assign the culture-based therapy. The proposed rule suggests that female gender and increasing age are key determinants for the potential of a patient to benefit from cycled therapy, which is not surprising since these characteristics are known risk factors for exacerbations. Counterintuitively, this rule suggests that *Pa* positivity at baseline leads against the recommendation for cycled therapy unlike with the prior outcome of increased *Pa* frequency. The rule also suggests that patients with genotypes other than *delf508* heterozygous and homozygous may benefit more from cycled therapy as compared to patient with the *delf508* mutation. Overall, these results indicate the importance of evaluating multiple endpoints to evaluate consistency in recommendations for treatment based on these proposed rules.

For RWL-VS-Gaussian, three clinical covariates (age, weight, baseline *S. aureus* status) were identified to be unimportant for estimating ITRs. It is well known that the shrinkage by the ℓ_1 penalty causes the estimates of the non-zero coefficients to be biased towards zero (Hastie et al. 2001). For similar reasons, the estimates from RWL-VS-Gaussian might be biased. One approach for reducing this bias is to run the ℓ_1 penalty method to identify important covariates, and then fit a model without the ℓ_1 term to the selected set of covariates. We applied RWL-Gaussian to the identified set of covariates. The results are also presented in Table 4. The obtained annual PE rate was better than that from RWL-Gaussian alone. Thus variable selection is still helpful to improve the performance.

In summary, the proposed RWL methods can identify potentially useful ITRs. For example, the estimated ITR for the *Pa* endpoint in the EPIC trial is that if the baseline *Pa* status is positive, assign culture-based therapy to the patient; if the baseline *Pa* status is negative, assign the cycled therapy. As the *Pa* endpoint was a secondary endpoint in the trial, the clinical utility of these findings must be weighed in context of all study findings and safety. But nonetheless identifying such a strategy is important to improve personalized clinical practice if the strategy is confirmed by future comparative studies.

7 Discussion

In this article, we have proposed a general framework called Residual Weighted Learning (RWL) to use outcome residuals to estimate optimal ITRs. The residuals may be obtained from linear models or generalized linear models, and hence RWL can handle various types of outcomes, including continuous, count and binary outcomes.

The proposed RWL methods appear to achieve better performance compared to other existing methods as shown in the simulation studies and real data analyses. The goal of RWL is to improve finite sample performance of OWL methods. It is not surprising that RWL outperforms OWL in terms of achieving higher value function and smaller variance. Two-step methods, such as ℓ_1 -PLS and ℓ_1 -PoissonReg, are generally parametric. Their performance largely depends on how close the posited model is to the true model of the conditional mean outcome. When the model is correctly specified, two-step methods can be more efficient than nonparametric RWL methods. However, the relationship among clinical

covariates, treatment assignment, and outcome is complex in practice. Misspecified models may lead to biased estimates.

Both two-step methods and RWL require correct model specification to guarantee their performance. We consider two levels of model specification. The first one is the model for the decision boundary; the other is the model for the outcome. Two-step methods, such as ℓ_1 -PLS, only specify models for conditional outcomes, and the decision boundary is implied from these models. Thus the performance of two-step methods are only related to these models. However, for RWL, the two levels are separated. RWL methods target the decision boundary directly. The model for decision boundary is more critical to performance. RWL with Gaussian RBF kernel can approximate any form of decision boundary when the sample size is sufficiently large. Even when using linear kernels, RWL is more robust to model misspecification on the decision boundary than the two-step methods. If the conditional mean outcome is assumed to be linear in the parameters, the decision boundary will be also; the converse does not hold. For RWL with linear kernel, the decision boundary remains a simple linear form, while the conditional mean outcome is allowed to be nonlinear.

Another source of robustness of RWL comes from the model for the outcome, in which we estimate residuals. By Theorem 3.3, no matter how bad the residual estimate is, RWL is still consistent. In the simulation studies, we have investigated the robustness of residual estimates on finite samples. We applied the null model, as an example of incorrect model specification, to estimate residuals, and achieved satisfactory performance. The simulation studies confirm the robustness of RWL. In this article, we used simple models to obtain the residuals. However, RWL is quite exible in the methods used to estimate residuals. When the sample size is small or the heterogeneous treatment effect is weak, we may consider a complicated method to accurately estimate residuals to improve the performance of RWL. Alternatively, nonparametric regression methods, such as support vector regression (Vapnik 1998) or random forests (Breiman 2001), are an option to eliminate model misspecification.

RWL is closely related to a robust augmented inverse probability weighted estimator (AIPWE) of the value function. AIPWE was first introduced into optimal treatment regimes by Zhang et al. (2012). Under the setup in this article, *i.e.*, $\mathcal{A} = \{1, -1\}$ and $\pi(a, \mathbf{x})$ is known, for any rule d , its associated value function can be estimated by an AIPWE,

$$AIPWE(d) = \frac{1}{n} \sum_{i=1}^n \left\{ \frac{\mathbb{I}(d(\mathbf{x}_i) = a_i)}{\pi(a_i, \mathbf{x}_i)} R_i - \frac{\mathbb{I}(d(\mathbf{x}_i) = a_i) - \pi(a_i, \mathbf{x}_i)}{\pi(a_i, \mathbf{x}_i)} \hat{m}(\mathbf{x}_i) \right\},$$

where $\hat{m}(\mathbf{x}) = \hat{\mu}_1(\mathbf{x}) \mathbb{I}(d(\mathbf{x}) = 1) + \hat{\mu}_{-1}(\mathbf{x}) \mathbb{I}(d(\mathbf{x}) = -1)$, $\hat{\mu}_1(\mathbf{x})$ is an estimate of $\mathbb{E}(R|\mathbf{X}=\mathbf{x}, A=1)$, and $\hat{\mu}_{-1}(\mathbf{x})$ is an estimate of $\mathbb{E}(R|\mathbf{X}=\mathbf{x}, A=-1)$. If we take $\hat{g}^*(\mathbf{x}) = \hat{m}(\mathbf{x})$, for a fixed $d(\mathbf{x})$, then maximizing AIPWE is equivalent to minimizing the risk in (6). Hence AIPWE is essentially a “contrast” weighted learning method, where the contrast is a modelled form that links closely to the residual. Conversely, if we take $\hat{\mu}_1(\mathbf{x}) = \hat{\mu}_{-1}(\mathbf{x}) = \hat{g}^*(\mathbf{x})$, where $\hat{g}^*(\mathbf{x})$ estimates $g^*(\mathbf{x})$ in (5), then maximizing AIPWE is again equivalent to minimizing (6). However, RWL can now be obtained by replacing the

0-1 loss with the smoothed ramp loss, and hence RWL can be viewed as a modification of AIPWE with alternative useful connections to classification methods.

Variable selection is critical for good performance of RWL, as demonstrated in the simulation studies and data analyses. We suggest that practical use of RWL should necessarily be accompanied by variable selection. In Section 4, we introduce the concepts of predictive and prescriptive covariates. It has been noticed in practical settings that the main effects, rather than treatment-covariate interactions, tend to explain most of the variability in the outcome, and thus they are important for good prediction (Gunter et al. 2011). It is not uncommon that interactions between prescriptive covariates and treatment are overlooked due to their small predictive ability. Hence variable selection techniques designed for prediction applications might not work perfectly in finding the prescriptive covariates. In contrast, RWL focuses on treatment-covariate interactions, and variable selection in RWL only targets the prescriptive covariates.

RWL methods with linear and Gaussian RBF kernels were discussed in this article. The decision rule from linear RWL is easy to interpret, but at the risk of misspecification. As a universal kernel, the Gaussian RBF kernel is free of model misspecification. An obvious downside is that the final decision rule may be difficult to interpret, and clinical practitioners may not be comfortable using “black box” decision rules. One compromise is to apply linear RWL with a rich set of bases, for example, including two-way and three-way covariate interactions.

When the endpoint outcome is the survival time, observations are commonly subject to right censoring because of subject dropout or censoring due to the end of study. Currently, RWL does not cover this type of outcomes in theory due to the censorship. We recommend using martingale residuals (Fleming and Harrington 1991) as alternative outcomes. Specifically, residuals in RWL could be opposites of martingale residuals that are estimated by a Cox proportional hazards model on clinical covariates but excluding treatment assignment. Actually, the idea is not new. Therneau et al. (1990) suggested that martingale residuals from a null Cox model could be used as outcomes for Classification and Regression Trees (CART). The rigorous theoretical justification is under development. Here we provide an intuitive explanation. Let T denote the survival time, and C denote the censoring time. The observed data quadruplet is $(\mathbf{X}, A, R=T \wedge C, \Delta=\mathbb{I}(T \leq C))$. Firstly, for general survival analysis, (R, Δ) is seen as the censored outcome. We view the survival data from a different perspective. We treat Δ as the outcome, and $(\mathbf{X}, R)$ as clinical covariates. In the framework of counting processes, the martingale residual can be explained as the difference between the observed number of failures (0 or 1) during the time in the study for a subject and the expected numbers based on the fitted model. Thus it is almost in line with our derivation of RWL, although R may not be appropriate as a clinical covariate. Secondly, martingale residuals from a Cox model fitted excluding a covariate of interest can be used to determine the appropriate functional form of the covariate (Therneau et al. 1990). If we estimate martingale residuals by ignoring treatment assignment, they contain information related to the heterogeneity of treatment effects. Thirdly, martingale residuals have some properties reminiscent of linear models, for example, the residuals are asymptotically uncorrelated.

Moreover, the weighted sum of martingale residuals is zero, which is essential for RWL to control the treatment matching factor.

Several extensions are worthy of further investigation. In this article, we only considered two-arm trials. In practice, some trials involve multiple arms, *i.e.* have three or more treatment arms. A recent review by Baron et al. (2013) reported that 17.6% of published randomized clinical trials in 2009 were multiple-arm. Because of the number of possible comparisons, it is more complex to find ITRs in such trials compared with two-arm trials. It would thus be worthwhile to extend RWL to multiple-arm trials. In the machine learning literature, two schools of approaches are generally suggested for multi-category classification. One is to directly take all classes into consideration. Multicategory SVMs (Lee et al. 2004) and penalized logistic regression (Zhu and Hastie 2004) are two examples. The other school is to construct and combine several binary classifiers. One-versus-one, one-versus-all, and more general error-correcting output codes (Dietterich and Bakiri 1995, ECOC) are in this school. Similar generalizations may be possible for finding ITRs for multiple-arm trials.

For chronic diseases, multi-stage dynamic treatment regimes are more useful than single-stage optimal treatment regimes. In the context of multi-stage decision problems, a dynamic treatment regime is a sequence of decision rules, one per stage of intervention, for adapting a treatment plan over time to an individual. Zhao et al. (2014) has extended OWL to dynamic treatment regimes. Considering the improved performance of RWL over OWL, it is of interest to extend RWL to dynamic treatment regimes.

The intended use of the estimated optimal ITRs discovered using the proposed approach is for treating new patients. The theoretical, simulation and data analysis results demonstrate that such estimated ITRs are likely to lead to improved clinical outcomes for these patients. However, since the data used for estimating the ITRs should not also be used for confirmation, it is expected that an additional randomized clinical trial comparing the estimated ITR with standard of care, or some other suitable treatment comparison, will be used for confirmation as is typically done for new candidate treatments (Zhao et al. 2011).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The EPIC trial was funded by the National Heart Lung and Blood Institute (NHLBI) and National Institute of Diabetes and Digestive and Kidney Diseases (NIDDK) (U01-HL080310) and the Cystic Fibrosis Foundation Therapeutics. The first and fourth authors were funded in part by the CTSA at the University of North Carolina at Chapel Hill (UL1TR001111) as well as by grant P01CA142538 from the National Cancer Institute (NCI). The second and third authors were funded in part by the CTSA at the University of Washington (UL1TR000423). The authors thank Editor Nicholas Jewell, associate editor and two reviewers for their helpful comments which led to a significantly improved paper.

References

Allen GI. Automatic Feature Selection via Weighted Kernels and Regularization. *Journal of Computational and Graphical Statistics*. 2013; 22(2):284–299.

- An LTH, Tao PD. Solving a Class of Linearly Constrained Indefinite Quadratic Problems by D.C. Algorithms. *Journal of Global Optimization*. 1997; 11(3):253–285.
- Argyriou, A., Hauser, R., Micchelli, CA., Pontil, M. In: Cohen, WW., Moore, A., editors. A DC-programming algorithm for kernel selection; ICML, volume 148 of ACM International Conference Proceeding Series; ACM. 2006; p. 41–48.
- Baron G, Perrodeau E, Boutron I, Ravaud P. Reporting of analyses from randomized controlled trials with multiple arms: a systematic review. *BMC Medicine*. 2013; 11(1):84. [PubMed: 23531230]
- Bartlett PL, Mendelson S. Rademacher and Gaussian Complexities: Risk Bounds and Structural Results. *Journal of Machine Learning Research*. 2002; 3:463–482.
- Breiman L. Random Forests. *Machine Learning*. 2001; 45(1):5–32.
- Byrd RH, Lu P, Nocedal J, Zhu C. A Limited Memory Algorithm for Bound Constrained Optimization. *SIAM Journal on Scientific and Statistical Computing*. 1995; 16(5):1190, 1208.
- Cohen, J. Statistical power analysis for the behavioral sciences. Lawrence Erlbaum Associates, Inc.; Hillsdale, NJ: 1988.
- Collobert, R., Sinz, F., Weston, J., Bottou, L. Trading convexity for scalability; Proceedings of the 23rd international conference on Machine learning, ICML '06, 201–208; New York, NY, USA: ACM. 2006;
- Devroye L, Györfi L, Lugosi G. A Probabilistic Theory of Pattern Recognition. Springer. 1996
- Dietterich TG, Bakiri G. Solving Multiclass Learning Problems via Error-correcting Output Codes. *Journal of Artificial Intelligence Research*. 1995; 2(1):263, 286.
- Fleming T, Harrington D. Counting Processes and Survival Analysis. Wiley. 1991
- Fung GM, Mangasarian OL. A Feature Selection Newton Method for Support Vector Machine Classification. *Computational Optimization and Applications*. 2004; 28(2):185–202.
- Grandvalet, Y., Canu, S. Adaptive Scaling for Feature Selection in SVMs. In: Becker, S., Thrun, S., Obermayer, K., editors. Advances in Neural Information Processing System. MIT Press; 2002. p. 553–560.
- Gunter L, Zhu J, Murphy S. Variable Selection for Qualitative Interactions. *Statistical Methodology*. 2011; 8:42–55.
- Hastie T, Tibshirani R, Friedman JH. The Elements of Statistical Learning. Springer. 2001
- Havlic DV, Kendall MA, Ive P, Kumwenda J, Swindells S, Qasba SS, Luetkemeyer AF, Hogg E, Rooney JF, Wu X, Hosseini pour MC, Laloo U, Veloso VG, Some FF, Kumarasamy N, Padayatchi N, Santos BR, Reid S, Hakim J, Mohapi L, Mugenyi P, Sanchez J, Lama JR, Pape JW, Sanchez A, Asmelash A, Moko E, Sawe F, Andersen J, Sanne I. Timing of Antiretroviral Therapy for HIV-1 Infection and Tuberculosis. *New England Journal of Medicine*. 2011; 365(16):1482, 1491. [PubMed: 22010914]
- Hsu, C-W., Chang, C-C., Lin, C-J. Technical report, Department of Computer Science. National Taiwan University, Taipei 106; Taiwan: 2003. A practical guide to support vector classification.
- Kimeldorf G, Wahba G. Some results on Tchebycheffian spline functions. *Journal of Mathematical analysis and applications*. 1971; 33:82, 95.
- Lafferty J, Wasserman L. Rodeo: Sparse, greedy nonparametric regression. *The Annals of Statistics*. 2008; 36(1):28–63.
- Lee Y, Lin Y, Wahba G. Multicategory Support Vector Machines, theory, and application to the classification of microarray data and satellite radiance data. *Journal of the American Statistical Association*. 2004; 99:67–81.
- Lin Y, Zhang HH. Component selection and smoothing in multivariate nonparametric regression. *The Annals of Statistics*. 2006; 34(5):2272–2297.
- Morales JL, Nocedal J. Remark on ‘Algorithm 778: L-BFGS-B: Fortran Subroutines for Large-scale Bound Constrained Optimization’. *ACM Transactions on Mathematical Software*. 2011; 38(1): 7:1–7:4.
- Murphy SA. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B*. 2003; 65(2):331–355.
- Murphy SA. An experimental design for the development of adaptive treatment strategies. *Statistics in medicine*. 2005; 24(10)

- Nocedal J. Updating Quasi-Newton Matrices with Limited Storage. *Mathematics of Computation*. 1980; 35(151):773–782.
- Nocedal, J., Wright, SJ. *Numerical Optimization*. 2nd edition. Springer; New York: 2006.
- Qian M, Murphy SA. Performance guarantees for individualized treatment rules. *The Annals of Statistics*. 2011; 39(2):1180–1210. [PubMed: 21666835]
- Rakotomamonjy A. Variable selection using svm based criteria. *J. Mach. Learn. Res.* 2003; 3:1357–1370.
- Robins, JM. In: Lin, D., Heagerty, P., editors. *Optimal Structural Nested Models for Optimal Sequential Decisions*; *Proceedings of the Second Seattle Symposium in Biostatistics*, volume 179 of *Lecture Notes in Statistics*, 189–326; Springer New York. 2004;
- Rudin, W. *Real and Complex Analysis*. 3rd Ed.. McGraw-Hill, Inc.; New York, NY, USA: 1987.
- Schmidt, M. *Graphical model structure learning with l1-regularization*. The University of British Columbia; 2010. Ph.D. thesis
- Steinwart I, Christmann A. *Support Vector Machines*. Springer. 2008
- Steinwart I, Scovel C. Fast rates for support vector machines using Gaussian kernels. *The Annals of Statistics*. 2007; 35(2):575–607.
- Therneau TM, Grambsch PM, Fleming TR. Martingale-based residuals for survival models. *Biometrika*. 1990; 77(1):147–160.
- Tibshirani R. Regression Shrinkage and Selection Via the Lasso. *Journal of the Royal Statistical Society, Series B*. 1994; 58:267–288.
- Treggiari MM, Retsch-Bogart G, Mayer-Hamblett N, Khan U, Kulich M, Kronmal R, Williams J, Hiatt P, Gibson RL, Spencer T, Orenstein D, Chatfield BA, Froh DK, Burns JL, Rosenfeld M, Ramsey BW. Comparative efficacy and safety of 4 randomized regimens to treat early pseudomonas aeruginosa infection in children with cystic fibrosis. *Archives of Pediatrics & Adolescent Medicine*. 2011; 165(9):847–856. [PubMed: 21893650]
- Treggiari MM, Rosenfeld M, Mayer-Hamblett N, Retsch-Bogart G, Gibson RL, Williams J, Emerson J, Kronmal RA, Ramsey BW. Early anti-pseudomonal acquisition in young patients with cystic fibrosis: Rationale and design of the EPIC clinical trial and observational study. *Contemporary Clinical Trials*. 2009; 30(3):256–268. [PubMed: 19470318]
- Vapnik, VN. *Statistical Learning Theory*. John Wiley and Sons; New York: 1998.
- Wang L, Zhu J, Zou H. Hybrid huberized support vector machines for microarray classification and gene selection. *Bioinformatics*. 2008; 24(3):412–419. [PubMed: 18175770]
- Weston, J., Mukherjee, S., Chapelle, O., Pontil, M., Poggio, T., Vapnik, V. *Advances in Neural Information Processing System*. MIT Press; 2000. Feature Selection for SVMs; p. 668–674.
- Wu Y, Liu Y. Robust truncated-hinge-loss support vector machines. *Journal of the American Statistical Association*. 2007; 102:974–983.
- Yuille AL, Rangarajan A. The concave-convex procedure. *Neural Computation*. 2003; 15(4):915–936. [PubMed: 12689392]
- Zhang B, Tsiatis AA, Laber EB, Davidian M. A Robust Method for Estimating Optimal Treatment Regimes. *Biometrics*. 2012; 68(4):1010–1018. [PubMed: 22550953]
- Zhao Y, Kosorok MR, Zeng D. Reinforcement learning design for cancer clinical trials. *Statistics in Medicine*. 2009; 28(26):3294–3315. [PubMed: 19750510]
- Zhao Y, Zeng D, Rush AJ, Kosorok MR. Estimating Individualized Treatment Rules Using Outcome Weighted Learning. *Journal of the American Statistical Association*. 2012; 107(499):1106–1118. [PubMed: 23630406]
- Zhao Y, Zeng D, Socinski MA, Kosorok MR. Reinforcement Learning Strategies for Clinical Trials in Nonsmall Cell Lung Cancer. *Biometrics*. 2011; 67(4):1422, 1433. [PubMed: 21385164]
- Zhao Y-Q, Zeng D, Laber EB, Kosorok MR. New Statistical Learning Methods for Estimating Optimal Dynamic Treatment Regimes. *Journal of the American Statistical Association*. 2014 In press.
- Zhu J, Hastie T. Classification of gene microarrays by penalized logistic regression. *Biostatistics*. 2004; 5(3):427–443. [PubMed: 15208204]
- Zhu, J., Rosset, S., Hastie, T., Tibshirani, R. *Advances in Neural Information Processing System*. 2003. 1-Norm Support Vector Machines.

Zou H, Hastie T. Regularization and variable selection via the Elastic Net. *Journal of the Royal Statistical Society, Series B*. 2005; 67:301–320.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Figure 1.

Example of residual weighted learning. The raw data are shown on the left, consisting of a single covariate X , treatment assignment $A = 1$ or -1 , and continuous outcome R with $E(R|X, A) = 3X + XA$. The allocation ratio is 1:1. The residuals, $R - 3X$, are shown on the right.

(a) Hinge loss

(b) Ramp loss

(c) Smoothed ramp loss

(d) Decomposition

Figure 2.

Hinge loss (a), ramp loss (b), and smoothed ramp loss (c) functions. (d) shows the difference of convex decomposition of the smoothed ramp loss, $T(u) = \phi_1(u) - \phi_0(u)$.

Table 1

Mean (std) of empirical value functions and misclassification rates evaluated on independent test data for 5 simulation scenarios with 5 covariates. The best value function and minimal misclassification rate for each scenario and sample size combination are in bold.

	<i>n</i> = 100		<i>n</i> = 400	
	Value	Misclassification	Value	Misclassification
Scenario 0 (Optimal value 1.43)				
ℓ^2 -PLS	1.41 (0.07)	0.05 (0.10)	1.42 (0.01)	0.02 (0.04)
OWL-Linear	1.39 (0.11)	0.07 (0.14)	1.43 (0.02)	0.01 (0.04)
OWL-Gaussian	1.33 (0.15)	0.13 (0.18)	1.41 (0.05)	0.03 (0.08)
RWL-Linear	1.40 (0.05)	0.06 (0.09)	1.42 (0.02)	0.03 (0.05)
RWL-Gaussian	1.34 (0.10)	0.13 (0.13)	1.40 (0.04)	0.07 (0.07)
Scenario 1 (Optimal value 2.25)				
ℓ^2 -PLS	2.22 (0.03)	0.06 (0.04)	2.24 (0.02)	0.03 (0.03)
OWL-Linear	1.91 (0.22)	0.23 (0.09)	2.08 (0.14)	0.16 (0.06)
OWL-Gaussian	1.88 (0.24)	0.24 (0.09)	2.08 (0.12)	0.16 (0.06)
RWL-Linear	2.19 (0.04)	0.09 (0.04)	2.23 (0.01)	0.05 (0.02)
RWL-Gaussian	2.17 (0.06)	0.11 (0.04)	2.22 (0.02)	0.05 (0.02)
Scenario 2 (Optimal value 1.96)				
ℓ^2 -PLS	1.71 (0.07)	0.24 (0.04)	1.75 (0.01)	0.22 (0.01)
OWL-Linear	1.51 (0.12)	0.32 (0.05)	1.59 (0.10)	0.29 (0.06)
OWL-Gaussian	1.49 (0.15)	0.33 (0.06)	1.63 (0.11)	0.26 (0.05)
RWL-Linear	1.66 (0.08)	0.26 (0.04)	1.74 (0.03)	0.23 (0.02)
RWL-Gaussian	1.75 (0.09)	0.20 (0.05)	1.90 (0.03)	0.10 (0.03)
Scenario 3 (Optimal value 3.88)				
ℓ^2 -PLS	3.00 (0.11)	0.38 (0.03)	3.03 (0.04)	0.37 (0.01)
OWL-Linear	2.98 (0.10)	0.38 (0.03)	3.01 (0.02)	0.38 (0.01)
OWL-Gaussian	3.21 (0.18)	0.32 (0.05)	3.56 (0.12)	0.22 (0.04)
RWL-Linear	3.15 (0.13)	0.34 (0.03)	3.28 (0.04)	0.31 (0.01)
RWL-Gaussian	3.62 (0.12)	0.19 (0.04)	3.82 (0.04)	0.10 (0.02)
Scenario 4 (Optimal value 3.87)				
ℓ^2 -PLS	2.29 (0.14)	0.49 (0.08)	2.38 (0.15)	0.55 (0.08)
OWL-Linear	2.42 (0.14)	0.55 (0.08)	2.49 (0.11)	0.59 (0.06)
OWL-Gaussian	2.43 (0.15)	0.52 (0.07)	2.59 (0.15)	0.50 (0.07)
RWL-Linear	2.42 (0.14)	0.54 (0.09)	2.49 (0.10)	0.58 (0.08)
RWL-Gaussian	2.68 (0.28)	0.43 (0.08)	3.49 (0.08)	0.22 (0.03)

Table 2

Mean (std) of empirical value functions and misclassification rates evaluated on independent test data for 5 simulation scenarios with 50 covariates. The best value function and minimal misclassification rate for each scenario and sample size combination are in bold.

	<i>n</i> = 100		<i>n</i> = 400	
	Value	Misclassification	Value	Misclassification
Scenario 0 (Optimal value 1.43)				
ℓ_1 -PLS	1.34 (0.15)	0.11 (0.17)	1.42 (0.03)	0.03 (0.06)
OWL-Linear	1.36 (0.15)	0.08 (0.17)	1.43 (0.03)	0.01 (0.03)
OWL-Gaussian	1.35 (0.14)	0.10 (0.16)	1.42 (0.04)	0.01 (0.05)
RWL-Linear	1.38 (0.11)	0.07 (0.13)	1.42 (0.04)	0.03 (0.06)
RWL-Gaussian	1.34 (0.13)	0.11 (0.15)	1.40 (0.05)	0.04 (0.07)
RWL-VS-Linear	1.34 (0.11)	0.11 (0.14)	1.40 (0.04)	0.05 (0.06)
RWL-VS-Gaussian	1.33 (0.12)	0.13 (0.14)	1.41 (0.03)	0.05 (0.06)
Scenario 1 (Optimal value 2.25)				
ℓ_1 -PLS	2.17 (0.11)	0.09 (0.06)	2.23 (0.02)	0.04 (0.03)
OWL-Linear	1.51 (0.12)	0.37 (0.03)	1.68 (0.10)	0.32 (0.03)
OWL-Gaussian	1.49 (0.13)	0.38 (0.04)	1.65 (0.11)	0.32 (0.04)
RWL-Linear	1.75 (0.14)	0.29 (0.05)	2.14 (0.03)	0.13 (0.02)
RWL-Gaussian	1.77 (0.12)	0.29 (0.04)	2.14 (0.03)	0.13 (0.02)
RWL-VS-Linear	2.05 (0.14)	0.17 (0.07)	2.22 (0.02)	0.06 (0.03)
RWL-VS-Gaussian	2.02 (0.18)	0.18 (0.08)	2.21 (0.08)	0.06 (0.04)
Scenario 2 (Optimal value 1.96)				
ℓ_1 -PLS	1.56 (0.19)	0.30 (0.07)	1.73 (0.03)	0.23 (0.01)
OWL-Linear	1.42 (0.14)	0.37 (0.04)	1.47 (0.06)	0.35 (0.01)
OWL-Gaussian	1.40 (0.14)	0.38 (0.04)	1.47 (0.07)	0.35 (0.02)
RWL-Linear	1.42 (0.12)	0.36 (0.04)	1.60 (0.05)	0.28 (0.02)
RWL-Gaussian	1.40 (0.14)	0.36 (0.04)	1.62 (0.05)	0.28 (0.02)
RWL-VS-Gaussian	1.56 (0.15)	0.30 (0.07)	1.92 (0.06)	0.09 (0.05)
Scenario 3 (Optimal value 3.88)				
ℓ_1 -PLS	2.90 (0.19)	0.40 (0.05)	3.00 (0.04)	0.38 (0.01)
OWL-Linear	2.96 (0.14)	0.39 (0.04)	3.00 (0.02)	0.38 (0.01)
OWL-Gaussian	2.97 (0.11)	0.39 (0.03)	3.00 (0.03)	0.38 (0.01)
RWL-Linear	2.90 (0.18)	0.40 (0.05)	3.00 (0.05)	0.38 (0.01)
RWL-Gaussian	2.88 (0.18)	0.41 (0.05)	2.96 (0.09)	0.39 (0.02)
RWL-VS-Gaussian	3.17 (0.33)	0.33 (0.09)	3.82 (0.13)	0.09 (0.05)
Scenario 4 (Optimal value 3.87)				

	<i>n</i> = 100		<i>n</i> = 400	
	Value	Misclassification	Value	Misclassification
$\hat{d}$ -PLS	2.32 (0.09)	0.49 (0.05)	2.36 (0.09)	0.52 (0.05)
OWL-Linear	2.41 (0.14)	0.55 (0.08)	2.48 (0.11)	0.59 (0.06)
OWL-Gaussian	2.43 (0.13)	0.55 (0.07)	2.47 (0.11)	0.58 (0.06)
RWL-Linear	2.38 (0.13)	0.53 (0.07)	2.44 (0.12)	0.56 (0.07)
RWL-Gaussian	2.37 (0.13)	0.53 (0.07)	2.44 (0.12)	0.56 (0.07)
RWL-VS-Gaussian	2.37 (0.12)	0.51 (0.05)	3.47 (0.47)	0.21 (0.15)

Table 3

Comparison of methods for estimating ITRs on the EPIC data with the Pa endpoint. Higher values are better.

	Predicted treatment #cycled vs #culture	Mean value (std)	Proportion of significant p-values	Median p-value
ℓ -PLS	168 : 115	0.913 (0.040)	1	0.004
OWL-Linear	283 : 0	0.894 (0.054)	0	0.111
OWL-Gaussian	283 : 0	0.894 (0.054)	0	0.111
RWL-Linear	178 : 105	0.896 (0.046)	0.37	0.072
RWL-Gaussian	169 : 114	0.900 (0.045)	0.64	0.0404
RWL-VS-Linear	168 : 115	0.909 (0.044)	0.90	0.002
RWL-VS-Gaussian	168 : 115	0.907 (0.045)	0.85	0.007
Fixed rule (cycled)	283 : 0	0.894 (0.054)	0	0.111
Fixed rule (culture-based)	0 : 283	0.864 (0.054)	0	0.889

Table 4

Comparison of methods for estimating ITR on the EPIC data with the PE endpoint. Lower annual PE rates are better.

	Predicted treatment #cycled vs #culture	Mean annual PE rate (std)	Proportion of significant p-values	Median p-value
PoissonReg	150 : 133	0.742 (0.224)	0.21	0.245
$\mathcal{L}$ -PoissonReg	120 : 163	0.707 (0.233)	0.55	0.042
OWL-Linear	141 : 142	0.722 (0.236)	0.27	0.123
OWL-Gaussian	140 : 143	0.714 (0.247)	0.33	0.079
RWL-Linear	105 : 178	0.673 (0.226)	0.82	0.013
RWL-Gaussian	108 : 175	0.667 (0.220)	0.90	0.008
RWL-VS-Linear	105 : 178	0.676 (0.227)	0.75	0.013
RWL-VS-Gaussian	108 : 175	0.688 (0.218)	0.69	0.024
RWL-VS-Gaussian + RWL-Gaussian	102 : 181	0.639 (0.220)	0.97	0.001
Fixed rule (cycled)	283 : 0	0.725 (0.245)	0	0.140
Fixed rule (culture-based)	0 : 283	0.823 (0.253)	0	0.860